Information Cascades with Noise

Tho Ngoc Le, Vijay G. Subramanian, and Randall A. Berry

Abstract

Online networks enable agents to better observe the behavior of others and in doing so potentially learn from their actions. A key feature of models for such social learning is that *information cascades* can result, in which agents ignore their private information and blindly follow the actions of other agents. This paper considers the impact of noise in the form of observation errors in such a model. Similar to a model without noise, we show that with noise, both correct and incorrect cascades happen, with the same level of fragility. However, with noise, it is harder to overturn a cascade from one direction to the other. Using Markov chain analysis, we derive the net welfare (payoff) of each agent as a function of his private signal quality and the error rate. We show, somewhat surprisingly, that in certain cases, increasing the observation error rate can lead to higher welfare for all but a finite number of agents. In such cases, we compare and contrast adding additional noise with simply withholding observations from the first few agents. Our analysis assumes that all erroneous observations are available on a common database; however, we also discuss relaxing this assumption. We conclude by discussing the impact of different tie-breaking rules on the probability of wrong cascade, and the impact of bounded rationality on the agents' welfare.

Index Terms

Perfect Bayesian Equilibrium, Bayesian Learning, Information Cascades, Herding.

I. Introduction

Consider a recommendation system where agents sequentially decide whether to buy an item, for which they have some prior knowledge of its quality/utility. Agents' decisions are reported to a common database that is available to all later agents (e.g., via a website), who can potentially

This work was supported by the NSF under grant IIS-1538827 and the Vietnam Education Foundation (VEF) Fellowship.
T. Le and R. Berry are with the EECS Dept of Northwestern University, and V. Subramanian is with the EECS Dept of the University of Michigan. Emails: thole2012@u.northwestern.edu, vgsubram@umich.edu, rberry@eecs.northwestern.edu. This paper was presented in part in [33], [34].

benefit from the information obtained by observing their predecessors' choices. The study of such a model was initiated in the seminal papers [9], [10], and [11], which cast this in an observational Bayesian learning framework. In these models, each individual (agent) has some prior knowledge or signal about some payoff-relevant state of the world generated according to a commonly known probability distribution. Each agent makes a one-time decision in an exogenous order and his action is recorded on a common database available to the public. From this database, each agent then observes exactly the decisions made by all previous agents. Given these observations and their own signals, agents are assumed to be Bayesian rational, i.e., they sequentially choose the action that reflects their posterior beliefs about the state of the world. To put this in game-theoretic jargon, the setting is equivalent to a dynamic game with asymmetric information. The equilibrium analyzed in these problems is a Perfect Bayesian Equilibrium (PBE), which uses the common information based belief to determine the strategies[1].

We highlight two possible outcomes for such a model. First, there is *learning*. This is defined as when the information in private signals is aggregated efficiently, and agents eventually know the underlying true value of the item and make socially optimal decisions. Second, there is *cascading*. An *information cascade* occurs when it is optimal for the agents to ignore their own signals when taking actions. Though individually optimal, this may lead to the agents making a choice that is socially sub-optimal, which we refer to as a "wrong cascade." In addition to the possibility of cascading to the wrong conclusion, an information cascade results in a loss of information about the private signals held by all the agents following the onset of cascade.

In models where a homogeneous population of agents with discrete, bounded private signals is assumed (e.g. [9]-[11]), a main result is that a cascade happens in a finite number of agents with probability one. This leads to a positive probability of cascading toward the wrong choice. Once a cascade occurs, all subsequent agents follow suit, given the underlying homogeneity. The ultimate consequence is that all private signals from the onset of cascading are lost; thus learning stops.

In this paper, we consider a similar Bayesian learning model as in [9]-[11] except we introduce *observation errors*. More precisely, we assume observation errors occur when each agent's action is recorded for all subsequent agents to see, with the recording subject to error, with the statistics

[1]See [25] where a formal proof is presented.

of the error process known to all agents. The motivation of having observation errors in our model comes from the following reasons. First, real-world agents can make mistakes. As an example, this could model a setting where agents are asked to report their decisions on a website and agents occasionally misreport (e.g., by simply pressing the wrong button). In communication networks settings, observation errors can also occur as a form of broadcast failure in feedforward networks ([29]). Indeed, even highly sophisticated decision makers such as investors ([28]) and consumers ([27]) can make mistakes in choosing (binary) actions (investing in stocks or purchasing goods). Second, observation errors could come from the mistakes incurred by learning algorithms based on web-log data. For example, consider scenarios when a website uses the customers' web-log records to infer the customers' purchasing decisions and then posts such decisions on a database; the errors can occur from both the data and the algorithms used for such inference (e.g., a user could buy an item from a different retailer or could buy it and then return it elsewhere). Finally, this observation error could also result from either strategic agents ([32]), or a social planner ([31]), who can manipulate the recording for their own benefit.[2] Even in scenarios when there is no clear incentive for benefit, agents can choose not to express their true opinions, e.g., IMDB review manipulations ([37]). We study the effects of these observation errors on the probabilities of cascades and the welfare obtained for any given agent, and also asymptotically as the number of agents increases without bound.

A. *Contributions*

Our work has several contributions. First, we develop a simple Markov chain model that we use to analyze the probabilities of cascades and the expected pay-off of each agent as a function of the error-rates and signal quality. We then study the effect of changing the noise level on both the asymptotic behavior and the behavior for an arbitrary agent by using stochastic ordering and coupling methods for Markov chains. Our results demonstrate a counter-intuitive phenomenon: for certain parameter settings, the asymptotic average payoff can be increased by increasing the noise level. The extent of this phenomenon and the amount of noise to be added depend on the agents' signal quality and the total amount of noise already present.

[2]In the case of strategic agents manipulating reviews, our notion of welfare would need to be modified to account for the preferences of these agents.

Second, we extend our conclusions by considering which agents benefit from adding noise in the low noise regime. Interestingly, we show that if one agent benefits, the following agent may not. However, we show several properties of the sequence of benefiting agents, and further show that if one agent benefits at most a finite number of subsequent agents will not benefit.

Third, we compare and contrast the trade-offs of two approaches for improving the agents' welfares: 1) Adding observation noise (whenever it is beneficial to do so); and 2) Forcing the first few agents to be "guinea pigs," i.e., to follow their private signals as in [19]. We see that the second method leads to higher asymptotic pay-offs, but at the cost of lower pay-offs for an initial group of users. We also suggest a hybrid method that for certain model parameters can significantly reduce the number of required guinea pigs, while still achieving better asymptotic welfare than either approach separately.

Fourth, we also discuss the role of a common database in our analysis. We argue that such a database is crucial in enabling the agents being able to to infer the exact time in the history where a cascade happens. In particular, we compare the common database setting to one where each agent's observation is an *i.i.d.* erroneous version of the sequence of the true actions taken before. In such case, we show that a cascade may not persist once it starts. Moreover, this is shown to greatly increase the complexity of updating beliefs for each agent, making the Bayesian assumption more questionable.

Finally, our work also highlights the advantage of Markov Chain methods in obtaining a refined and detail analysis of the problem, in particular, obtaining per agent payoffs and determining sensitivity to parameter changes. This should be contrasted with the most of the literature, which uses martingale methods to discuss asymptotic performance measures.

The remainder of the paper is organized as follows. We start by discussing related work in Section I-B. In Section II, we state our model. In Section III, we analyze the Bayesian updates for this model and give our Markov chain formulation. We then turn to studying the impact of changing the noise level in Section IV. In Section V, our analysis of when adding noise is better is presented, for both cases of asymptotic and individual welfares. We compare adding noise to the "guinea pigs" model in Section VI, and then consider the case of no common database in Section VII. In Section VIII, we present some discussions of other changes to the model including different tie-breaking rules, and the effect of agents' bounded rationality. We conclude in Section IX.

B. Related work

In models where information is aggregated independently and efficiently, e.g. [22], learning can be achieved. However, the results in [9]-[11] showed that pathological outcomes, such as cascading/herding can occur in scenarios where decisions are taken sequentially and agents can observe previous actions. In the following, we organize the related literature in an attempt to answer the questions: why does a cascade happen? or when can learning occur?

First, learning fails because the discrete feedback, i.e. previous actions, from the agents is not sufficiently rich. This is connected to early models of sequential detection/hypothesis testing, e.g., [2] and [3]. In these works, a hypothesis is tested at each stage based on a past feedback and a new observation. A new feedback is then encoded and passed on to the next stage. In [3], the authors showed that if a discrete feedback is used with a regular encoding-decoding procedure, then the probability of choosing the wrong hypothesis is bounded away from zero. In terms of cascading jargon as in [9]-[11], this means the probability of wrong cascade is positive. Instead, by changing the encoding-decoding algorithm, the authors in [2] showed that learning is achieved using a feedback encoded in only two bits of information.

Second, as showed by Smith and Sorensen in [17], the boundedness of the likelihood ratios of the private signals is another reason why learning with the regular encoding-decoding procedure fails. Using both martingale techniques and Markovian analysis, Smith and Sorensen showed that with unbounded likelihood ratio, any deep cascade can be stopped and learning toward the correct action can be achieved. Otherwise, there is a positive probability that learning fails and all but a first few agents cascade to the less profitable action. In addition, these authors considered a richer assumption where there are two types of agents with the types unknown. Assuming that the distribution of the types is common knowledge, this leads to another type of outcome known as "confounded learning", where from some point onward the history offers no information and subsequent agents must rely on their private signals.

Third, another reason is that Bayesian update plus constant threshold rule may not be optimal to induce learning. In [2], the author proposed a time-variant encoding rule for finite memory approximation that make learning possible. Similar conclusions also hold for more general networks (e.g., [29]), or for a random decision rule (e.g., [35]). In these works, it is not clear whether it is in the agents' best interest to strictly adhere to such set of rules; therefore the

outcomes might not constitute an equilibrium. Here, we instead consider a strategic setting where each agent acts in his own benefit, and the agents are in a PBE. In particular, we limit ourselves to binary feedback in the form of the agent's actions and assume that each agent deterministically chooses his action to optimize his own pay-off (i.e., there is no planner).

Finally, the richest venue for inducing learning, as being explored by the literature, is by changing the *information structure* from the simple models in [9]-[11]. In the rest of this section, we discuss works that focus on this approach. In [25], the authors also considered Bayesian, observational learning with bounded and unbounded signal quality (similar to [17]). However, each agent only gets to see the actions of a random subset (with some underlying distribution) of the agents that acted before. In particular, the authors provided the conditions on the network formed by sampling past actions that would support learning. If private signals are unbounded, learning is achieved if such network satisfies "expanding observation"[3]. Moreover, even when the private signals are bounded, if the network allows information to be collected and passed to later agents, learning is also achieved.

Another work on observational learning (with both Bayesian and non-Bayesian settings) when the signals are bounded, and the network is also sampled, was presented in [30]. However, the sampling scheme is deterministic as the authors assumed that each agent observes only a fixed K of his immediate predecessors. In the Bayesian setting where the agents myopically maximize their expected payoffs, learning almost surely fails for all K while it is achieved in probability only for $K \geq 2$ by using a sophisticated encoding-decoding procedure, as in the spirit of [2]. On the other hand, for the non-Bayesian setting where agents are forward looking and maximize the discounted sum (over all agents) of the probabilities of a correct decision, learning almost surely fails for all values of K.

Also in the line of sampling of the observation history, in [36] the author instead studied a social learning model where agents move in endogenous order. In this setting, the agents have to pay a cost to, strategically, choose their observation networks. Under these assumptions, the model showed that unbounded private belief is not necessary for asymptotic learning as long as

[3]According to these authors, an expanding network means agents avoid observing exclusively any subset of other agent, thus allow the arrival of new information from other sources in the network.

agent are not limited to seeing only a finite number of observations[4]. Similar to [25], the author characterized the pure-strategy PBE for which learning occurs.

In another strand of work where the information structure is modified, an attempt to improve learning that aims to achieve maximum social-welfare is by sacrificing a number of agents as per the sequential experiment policy for the multi-armed bandit problem suggested by [7], [13]. Based on such experiments, we can more accurately solving the hypothesis test problem and use the result for remaining agents. An example in this line of work was suggested by Sgroi in [19], where the sacrificial victims are called the "guinea pigs".

In Sgroi's paper [19], the author considered a similar model to [9], assuming a population with finite number of agents. Suppose that there is a social planner who can force the first few agents, the "guinea pigs," to make decisions based only on their private signals. As a result, subsequent agents, who benefited from collecting the additional private signals from those guinea pigs, could make a decision that leads to higher expected welfare. Thus, as more guinea pigs are sacrificed, more useful private signals are collected, and a closer to the socially optimal decision is achieved in the long run. The author also found the optimal number of guinea pigs that would maximize the total social welfare in a finite population. From a mechanism design viewpoint, however, the model in [19] omitted a discussion on the selection criteria of the potential guinea pigs and their corresponding fairness measures. Moreover, in a population with n agents, the optimal number of guinea pigs needed is $O(\log n)$ which increases without bound ([13]).

Another modification to the information structure includes learning with continuous action spaces. Given that every agent is endowed an informative private signal about the state of the world, learning could be achieved if these signals are appropriately collected. Pathologically, the occurrence of herding shows how such information is poorly aggregated ([15], [17], [12]). In [12], the author showed that herding can be prevented if there exists an action set that is a one-to-one mapping with the posterior of each agent. In other words, by allowing a continuum action space, each agent can fine-tune his action to the posterior. However, in practice, the assumption that every action can be chosen from a uniformly fine-grained set is rather strict for most decision making scenarios. Our paper and most of other works on the topic, instead, assume that the action space is discrete.

[4]This satisfies the "expanding network" condition as defined in [25]

In this paper, we assume a bounded signal quality, where types of agents are homogeneous, and a simple sampling network with all past actions being observable. However, our paper considers a different approach toward modifying the information structure. In addition to the asymptotic results, we also study finer per agent characterization and the variational characterization (as parameters change) that the Markovian analysis and coupling techniques afford. In particular, we relax the assumption of perfect observability of the action history and consider a similar Bayesian learning model as in [9]-[11] except with the introduction of noise in the form of *observation errors*. We study the effects of such errors on the welfare obtained asymptotically as the number of agents increases without bound. More precisely, we assume observation errors occur when each agent's actions are recorded on a common database, but this record is subject to errors, with the statistics of the errors process known to all agents.

II. Model

We consider a model similar to [9] in which there is a countable population of agents, indexed $n = 1, 2, \ldots$ with the index reflecting both the time and the order in which agents act. Each agent n chooses an action A_n of either buying (Y) or not buying (N) a new item. The true value (V) of the item can be either good (G) or bad (B); for simplicity, both possibilities are assumed to be equally likely.[5]

The agents are Bayes-rational utility maximizers[6] whose payoff structure is based on the agent's action and the true value V. If an agent chooses N, his payoff is 0. On the other hand, if he chooses Y, he faces a cost of $C = 1/2$ and gains one of two amounts depending on the true value of the item: his gain is 0 is $V = B$ and 1 if $V = G$. The total pay-off of an agent choosing Y is then the gain minus the cost. Thus, the *ex ante* expected payoff of each agent is 0.

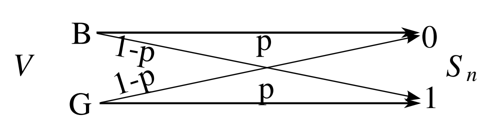

Figure 1: The BSC over which agents receive signals.

[5]A discussion of how to generalize this is provided in Section VIII-C .

[6]This assumption leads to a PBE that uses the common information based belief to determine strategies as proved in [25].

To reflect the agents' prior knowledge about V, each agent n receives a private signal $S_n \in \{1\text{ (high)}, 0\text{ (low)}\}$ through a binary symmetric channel (BSC) with crossover probability $1-p$, where $0.5 < p < 1$. (See Fig. 1.) In other words, we have

$$\mathbb{P}(S_n = 1|V = G) = \mathbb{P}(S_n = 0|V = B) = p, \text{ and}$$
$$\mathbb{P}(S_n = 0|V = G) = \mathbb{P}(S_n = 1|V = B) = 1 - p. \tag{1}$$

Thus, the private signals are informative, but not revealing. We modify the information structure in [9] by assuming that later agents' observations are noisy versions of their predecessors' actions. Specifically, we assume that each agent reports his action to a public database which is available to all successors. The errors in this process are modeled by passing every action A_i through another BSC with crossover probability $\epsilon \in (0, 1/2)$. This means with probability $1-\epsilon$, the reported action of agent i, O_i, satisfies $O_i = A_i$, and with probability ϵ, $O_i = \overline{A_i}$, where $\overline{A_i}$ is the opposite action of A_i. This assumption reduces the dependence of every agent's decision on the predecessors' choices and drives him toward using his own signal. With this in place we assume that each agent n takes a one-time action A_n based on his private signal S_n and the (noisy) observations $O_1, \ldots, O_{n-1}$ of all previous agents' actions $A_1, \ldots, A_{n-1}$.

III. Bayesian updates, cascades and error thresholds

A. Bayesian updates

The optimal action for the first agent is always to follow his private signal since no observation history is available. Starting from the second agent, every agent n considers his private signal S_n and the observations $O_1, ..., O_{n-1}$. Let the information set of agent n be $\{S_n, \mathcal{H}_{n-1}\}$ where $\mathcal{H}_{n-1} = \{O_1, ..., O_{n-1}\}$ is the observation history. Based on $\{S_n, \mathcal{H}_{n-1}\}$, agent n will update his posterior probability (using Bayes formula).

Definition 1. *The posterior probability of agent n, γ_n, is defined as* $\gamma_n(S_n, \mathcal{H}_{n-1}) = \mathbb{P}[V = G|S_n, \mathcal{H}_{n-1}]$.

Assuming common rationality[7], the decision rule for agent n is as follows:[8]

$$A_n = \begin{cases} Y, & \text{if } \gamma_n > 1/2, \\ N, & \text{if } \gamma_n < 1/2, \\ \text{follows } S_n, & \text{if } \gamma_n = 1/2. \end{cases} \tag{2}$$

Following [17], we can also state this decision rule in terms of a private and public likelihood ratio defined next.

Definition 2. *The private likelihood ratio of agent n, β_n, is defined as $\beta_n(S_n) = \mathbb{P}[S_n|V = B]/\mathbb{P}[S_n|V = G]$.*

Note that the private likelihood ratio is a function of the private signal, and we have $\beta_n(0) = p/(1-p), \beta_n(1) = (1-p)/p$. For a given $p \in (0.5, 1)$, these private ratios are strictly bounded in $(1, \infty)$ and $(0, 1)$, respectively.

Definition 3. *The public likelihood ratio after agent n decides, ℓ_n, is defined as $\ell_n(\mathcal{H}_n) = \mathbb{P}[\mathcal{H}_n|V = B]/\mathbb{P}[\mathcal{H}_n|V = G]$. This is available to all subsequent agents $n+1, n+2, \ldots$*

A simple application of Bayes formula gives $\gamma_n = \frac{1}{1+\beta_n \ell_{n-1}}$. Therefore, using the above formulae for β_n and ℓ_{n-1}, the rule in (2) can be applied to the resulting γ_n, which amounts to comparing $\beta_n \ell_{n-1}$ to 1, e.g., with Y being the choice whenever $\beta_n \ell_{n-1} < 1$. In this decision-making process, cascading is defined as follows:

Definition 4. *An information cascade is said to happen when an agent chooses some fixed action regardless of his private signal.*

In particular, agent n cascades to the action $Y(N)$ if $\gamma_n > 1/2$ $(< 1/2)$ for all realizations of $S_n \in \{0, 1\}$. In other words, agent n cascades to Y if $\beta_n(0)\ell_{n-1} < 1$ (i.e., $\ell_{n-1} < (1-p)/p$)), and cascades to N if $\beta_n(1)\ell_{n-1} > 1$ (or $\ell_{n-1} > p/(1-p)$). It was shown in [9] that when agents are rational and observations are perfect (i.e., $O_n = A_n, \forall n = 1, 2, \ldots$), an information cascade occurs in finite time with probability one. In our paper, we will show that this still holds. Another related

[7]As in the PBE in dynamic games.

[8]When equality holds, our decision rule differs from [9], where it is assumed that indifferent agents randomly choose one action. Our assumption not only simplifies the analysis but also proves to be the best tie-breaking rule (see Section VIII-A).

phenomenon, *herding*, happens when agents choose to copy the actions of others irrespective of the realizations of their signals. In a model with a homogeneous population and discrete, bounded private signals (e.g. [9]-[11] and our paper), the two concepts cascading and herding are equivalent. Thus, from now on we will use the two interchangeably.

B. Cascading properties

In this section, we outline some basic properties of cascading with observation errors. These naturally extend properties for the noiseless case shown in [9], [19], so we omit detailed derivations. For completeness we provide complete proofs in Section VIII-C.

Property 1. *Until cascading occurs, agent n's Bayesian updates depend only on his private signal S_n, and the difference in the numbers of Y's and N's in the observation history $\mathcal{H}_{n-1}$.*

In other words, given the common observation database $\mathcal{H}_{n-1}$, this difference of Y's and N's in $\mathcal{H}_{n-1}$ is a sufficient statistic for the observation history of agent n. This follows from the symmetry of the signal quality and the errors, which enables opposite observations to be "cancelled."[9]

Note that once a cascade happens, any new observations do not provide new information about the true value V, thus the public likelihood ratio stays the same. On the other hand, if an agent n does not cascade, all agents $1, 2, \ldots, n$ act independently based on their own private signals. Therefore, until a cascade happens, the public likelihood ratio ℓ_n depends only on the difference in the numbers of Y's and N's in the observation history $\mathcal{H}_n$. Specifically, let h_n be the number of Y's minus the number of N's in $\mathcal{H}_n$ (prior to a cascade), then $\ell_n = \left(\frac{1-a}{a}\right)^{h_n}$, where for all $i = 1, 2, \ldots, n$:

$$a = \mathbb{P}[O_i = Y|V = G] = \mathbb{P}[O_i = N|V = B], \tag{3}$$

gives the probability that a given observation agrees with the optimal action for an agent that is not cascading. We can show that $a = f(\epsilon, p) \in (0.5, p)$, where:

$$f(x, y) \triangleq x(1 - y) + (1 - x)y. \tag{4}$$

[9]This also extends to the case when V is not equally likely G or B as shown in Section VIII-C.

Property 2. *Once a cascade happens, it lasts forever.*

Once started, a cascade makes agents stop using their private signals and and so they provide no information to the successors. The successors are left in the same situation with the optimal choice of the first one who started the herd. We stress the importance of having the common database available to all agents. With such information, every agent updates h_n and knows the exact time when a cascade happens. All other subsequent observations from the onset of a cascade will be ignored. Thus, there is no need to further update h_n. Next, we provide a condition for determining the onset of cascades given the signal quality p and observation error ϵ.

C. Error thresholds

Without noise ($\epsilon = 0$), it was shown in [9] that cascading starts with the first agent n to have $|h_{n-1}| = 2$, i.e., the first agent to observe one action two times more than the other. However, as ϵ increases, each observation provides less information, which in turn can increase the value of $|h_n|$ needed to start a cascade. This is characterized in the following lemma.

Lemma 1. *Let $\alpha = (1-p)/p$. For any $k \geq 2$, define the increasing sequence of thresholds $\{\epsilon_k^*\}_{k=2}^{\infty}$, where the k^{th} threshold ϵ_k^* is given as:*

$$\epsilon_k^* = \left(1 - \alpha^{\frac{k-2}{k-1}}\right) / \left(1 - \alpha^{\frac{k-2}{k-1}} + \alpha^{\frac{-1}{k-1}} - \alpha\right). \tag{5}$$

Define $\mathcal{I}_k \triangleq \left[\epsilon_k^*, \epsilon_{k+1}^*\right)$ as the k^{th} ϵ–interval. Then for $\epsilon \in \mathcal{I}_k$, any agent n starts a cascade as soon as $|h_{n-1}| = k$.

The proof follows from direct calculation of γ_n; a detailed version is given in the Appendix. Note also that the parameter α here is the same as $\beta_n(1)$. Fig. 2 shows the thresholds ϵ_k^* for different values of k and p. For $k = 2$, this lemma yields $\epsilon_2^* = 0$, which is the case of noiseless observations as in [9].

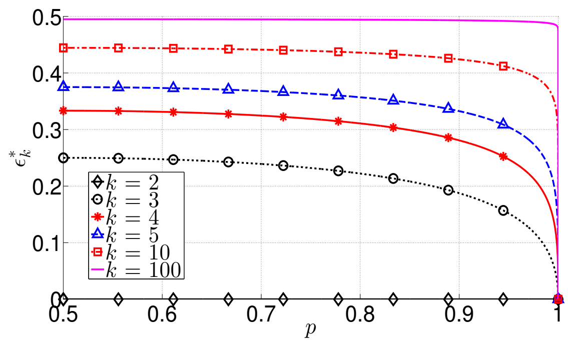

Figure 2: Thresholds for ϵ

A direct consequence of this lemma is that we can deduce k from a given ϵ as:

$$k = \left\lfloor \log_{(1-a)/a} \alpha \right\rfloor + 1. \tag{6}$$

Here, recall that a is a function of ϵ. Thus, given the signal quality p, the observation error $\epsilon \in \mathcal{I}_k$ and the observation history $\{O_1, O_2, \ldots, O_{n-1}\}$, an agent n can track the difference h_{n-1} and declare a N cascade as soon as $h_{n-1} = -k$, and a Y cascade as soon as $h_{n-1} = k$. In the limit of $\epsilon \to 0.5$, we have $k \to \infty$, and thus agents only use their own signals and herding never starts. In this case, no information is passed through each observation. Therefore, no learning occurs either.

Also note that ϵ_k^* is a decreasing function of p. Therefore, the more accurate the private signal, the sooner cascades may start for a fixed ϵ. In particular, Fig. 2 shows that if the private signal quality is high, with an intermediate level of noise a cascade may occur very early.

In the models without error ([9]-[11]), the authors also discussed how *fragile* a cascade can be. In fact, for the noiseless case, an information cascade occurs merely on the agreement of the running count of an action exceeding the other by two. Assuming that any of the subsequent agents is endowed with an additional private signal that disagrees with the cascading actions, the cascade stops. This provides some intuition to why many asset markets are volatile, or why changes in fashions happen frequently and unpredictably. With noisy observations, our results show that information cascades have the same level of fragility as in the noiseless case. In fact, for $\epsilon \in \mathcal{I}_k$, it would require only one additional opposite private signal to reduce the difference to $|h_{n-1}| < k$, so that the cascade stops. However, with noise, k can exceed 2 and it is harder to

reverse a cascade to the opposite direction: one would need a longer sequence of agents with private signals indicating the other action choice. Therefore, given a realistic setting where the information is noisy, it may take a long time to reverse certain trends.

D. Markovian analysis of cascades

By symmetry, first consider the case $V = G$.[10] By Property 1, a non-cascading agent n's observation history can be summarized by h_n. Thus, viewing each agent as a time-epoch, the agent's observation history can be represented as a state of a discrete-time Markov chain (DTMC). Each state i represents values of h_{n-1} that agent n may see before making his decision. Note that the first agent starts at state 0 since no observation history is available.

For the rest of the paper, assume $\epsilon \in \mathcal{I}_k$ so that an agent n will start a cascade if and only if $|h_{n-1}| = k$. Thus, to model cascades, we simply make the states $\pm k$ absorbing, and so are left with a finite state DTMC. The two events N cascade and Y cascade translate into hitting the left ($-k$) and the right (k) walls (or absorbing states), respectively. The probability of moving one step to the right is the probability that one more Y is added to the observation history, i.e., $a = \mathbb{P}[O_n = Y|V = G] = f(\epsilon, p) > 0.5$. Likewise, the probability of moving one-step to the left is $1 - a$. Hence, this DTMC is a simple random walk as shown in Fig. 3.

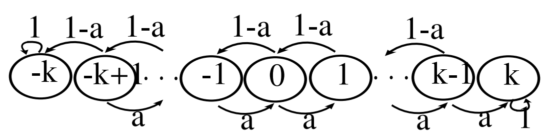

Figure 3: Transition diagram of the random walk when $V = G$.

The state transition matrix of this DTMC is given by

$$Q = \begin{bmatrix} 1 & 0 & 0 & ... & 0 & 0 & 0 \\ 1-a & 0 & a & ... & 0 & 0 & 0 \\ . & . & . & . & . & . & . \\ 0 & 0 & 0 & ... & 1-a & 0 & a \\ 0 & 0 & 0 & ... & 0 & 0 & 1 \end{bmatrix}. \tag{7}$$

[10]If V is not equally likely G or B, the dynamics of the Markov chains for this analysis are unchanged since these are conditioned on the true V. Moreover, the (unconditional) probabilities of cascades and expected welfare can be easily adjusted according to the prior knowledge of V. See details in Section VIII-C.

Since $a > 0.5$, Q is a row stochastic matrix with a drift to the right. We will use methods in [1] to calculate the probability of being at either absorbing state at any given time; since the agent indices correspond to the time index, this yields the probability of cascade for each agent. Assume that the process starts at state i. Let $u^*_{i,n}, v^*_{i,n}$ be the probabilities of being at the left wall, $-k$, and the right wall, k, at the n^{th} step, respectively. Let $u_{i,n}, v_{i,n}$ be the probabilities of hitting the left wall and the right wall for the first time at the n^{th} step, respectively. Obviously, $u^*_{0,n} = v^*_{0,n} = 0$ for $1 \leq n \leq k-1$. By [33], $u_{0,n} = v_{0,n} = 0$, for $n-k$ odd, thus for $n \geq k$:

$$u^*_{0,n} = \sum\nolimits_{m=k}^{n} u^*_{-k,n-m} u_{0,m} = \sum\nolimits_{m=k}^{n} u_{0,m}, \tag{8}$$

$$v^*_{0,n} = \sum\nolimits_{m=k}^{n} v^*_{k,n-m} v_{0,m} = \sum\nolimits_{m=k}^{n} v_{0,m}, \tag{9}$$

where the explicit expressions for the terms on the right-hand side in (8) and (9) are given in the next lemma. The proof can be found in Lemma 2 of [33].

Lemma 2. *For $n-k$ even:*

$$u_{0,n} = (1/k)2^n a^{\frac{n-k}{2}} (1-a)^{\frac{n+k}{2}} A_{k,n}, \tag{10}$$

$$v_{0,n} = (1/k)2^n a^{\frac{n+k}{2}} (1-a)^{\frac{n-k}{2}} A_{k,n}, \tag{11}$$

where

$$A_{k,n} = \sum\nolimits_{j=1,\, odd}^{j<k} \cos^{n-1}\left[j\pi/(2k)\right] \sin\left[j\pi/(2k)\right] (-1)^{\frac{j-1}{2}}. \tag{12}$$

By symmetry, if $V = B$, the wrong and correct cascade probabilities for agent n are $v^*_{0,n}$ and $u^*_{0,n}$, respectively. The following lemma states some useful relations between the probabilities of cascade.

Lemma 3. *For $\epsilon \in \mathcal{I}_k$, and any arbitrary agent $n \geq k$:*

1) If k is even, then $u^*_{0,n} = u^*_{0,n+1}, v^*_{0,n} = v^*_{0,n+1}$ for n even.

2) If k is odd, then $u^*_{0,n} = u^*_{0,n+1}, v^*_{0,n} = v^*_{0,n+1}$ for n odd.

3) The probabilities of correct and wrong cascade are related by:

$$v^*_{0,n}/u^*_{0,n} = [a/(1-a)]^k. \tag{13}$$

The proof of part 1) and 2) follows directly from (8) and (9), whereas the proof for part 3) follows by using the formulae from Lemma 2 in (8) and (9).

IV. Effect of varying observation error rates

We next turn to studying the effect of the error rates on the cascade probabilities and agent welfare.

A. Cascade probabilities

Lemma 2 shows that the probabilities of wrong and correct cascade depend on k and a. From Lemma 1 and since $a = f(\epsilon, p)$, it follows that for a fixed p, these probabilities are determined by the error probability ϵ.

The next theorem characterizes the effect of varying the error ϵ on the probabilities of wrong cascade, $u^*_{0,n}$, and correct cascade, $v^*_{0,n}$, for an arbitrary agent n.

Theorem 1. *For $\epsilon \in \mathcal{I}_k$:*

1) *The wrong cascade probability increases with ϵ.*
2) *The correct cascade probability decreases with ϵ.*

Proof. Consider $\epsilon^{'} < \epsilon^{''}$, both in $\mathcal{I}_k$, and let $\{Z^{'}_n\}_{n\geq 0}$ and $\{Z^{''}_n\}_{n\geq 0}$ be the two corresponding DTMCs on the same state space $\mathcal{S} = \{-k, -k+1, ..., 0, ..., k-1, k\}$. The following concept of *stochastic ordering* ([6]) compares these two DTMCs:

Definition 5. *Let X and Y be two discrete random variables taking values on the same set $\mathcal{S}$ and let $\boldsymbol{x}$ and $\boldsymbol{y}$ be their corresponding probability distribution vectors. $X(\boldsymbol{x})$ is said to be larger than $Y(\boldsymbol{y})$ in stochastic ordering, denoted by $X \geq_{st} Y(\boldsymbol{x} \geq_{st} \boldsymbol{y})$, if*

$$\sum\nolimits_{i\geq j} x_i \geq \sum\nolimits_{i\geq j} y_i, \text{ for all } j \in \mathcal{S}. \tag{14}$$

Definition 6. *The DTMC $\{Z^{'}_n\}$ is said to be larger than the DTMC $\{Z^{''}_n\}$ in stochastic ordering, denoted by $\{Z^{'}_n\} \geq_{st} \{Z^{''}_n\}$, if*

$$Z^{'}_n \geq_{st} Z^{''}_n, \text{ for all } n \geq 0. \tag{15}$$

Definition 7. *A transition matrix Q is said to be stochastically increasing if for all $i, i-1 \in \mathcal{S}$:*

$$\boldsymbol{Q_i} \geq_{st} \boldsymbol{Q_{i-1}} \tag{16}$$

where $\boldsymbol{Q_i}$ denotes the i^{th} row of Q (which is a distribution vector for some random variable).

The proof continues by noting that the corresponding right transition probabilities of the two DTMC's satisfy $a^{'} > a^{''} > 0.5$. Thus, $\boldsymbol{Q^{'}_i} \geq_{st} \boldsymbol{Q^{''}_i}$ for all $i \in \mathcal{S}$. Moreover, the transition matrices

for each DTMC are stochastically increasing, and both $\{Z_n'\}$ and $\{Z_n''\}$ start from the same state 0. Therefore, by Theorem 4.2.5a and equation (4.2.16) in [6], we have $\{Z_n'\} \geq_{st} \{Z_n''\}$. By Defn. 6, it follows that at any arbitrary time n, $Z_n' \geq_{st} Z_n''$. Let z' and z'' be the corresponding probability distribution vectors at time n.

Setting $j = k$ in Defn. 5, we have $v_{0,n}^{*'} = z_k' \geq z_k'' = v_{0,n}^{*''}$. Now, letting $j = -k + 1$ yields $\sum_{i \geq -k+1} z_i' \geq \sum_{i \geq -k+1} z_i''$, so that $u_{0,n}^{*'} \leq u_{0,n}^{*''}$. If fact, we can use a coupling argument to prove that equality does not hold for $n > k$, i.e., $v_{0,n}^{*'}$ is strictly greater than $v_{0,n}^{*''}$ for $n > k$. These details are shown in the Appendix E. Similarly, $u_{0,n}^{*''} > u_{0,n}^{*'}$. Since Lemma 2 shows that $u_{0,n}^{*}$ and $v_{0,n}^{*}$ are continuous functions of ϵ (in the specified interval), this completes the proof. $\square$

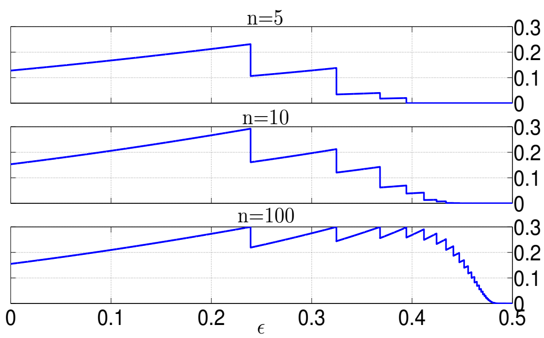

Figure 4: Probability of wrong cascade for agent n, with $p = 0.70$.

As a demonstration of Theorem 1, Fig. 4 shows the probability of wrong cascade when $p = 0.70$, for agents $n = 5, 10$ and 100. Notice that for any agent n, Lemma 1 shows that there exists ϵ close enough to 0.5 that yields $k \geq n$. The probability of wrong cascade will thus go to zero for large enough ϵ, as shown in the figure. Note also that for each agent n, as ϵ increases, the probability of wrong cascade discontinuously decreases at a countable number of points (and continuously increases in between). These points correspond exactly to the values ϵ_k^* for different choices of integer $k \geq 2$.

B. Agent Welfare

Let π_n be the payoff or welfare of agent n. From Section II, we have $\pi_n = 0$ if $A_n = N$, while if $A_n = Y$, π_n is either $1/2$ or $-1/2$ corresponding to $V = G$ or $V = B$, respectively. All agents $1 \leq n \leq k$ use their own signals; thus by averaging over $V \in \{G, B\}$ they all have the same

welfare given by:

$$E[\pi_n] = \{\mathbb{P}[A_n = Y|V = G] - \mathbb{P}[A_n = Y|V = B]\}/4$$
$$= (2p-1)/4 \triangleq F. \tag{17}$$

For agents $n \geq k+1$:

$$E[\pi_n] = \{\mathbb{P}[A_n = Y|V = G] - \mathbb{P}[A_n = Y|V = B]\}/4$$
$$= F + \left[(1-p)v^*_{0,n-1} - pu^*_{0,n-1}\right]/2. \tag{18}$$

For a fixed observation error ϵ, (18) explicitly relates the welfare of an agent n, who faces the possibility of cascade, to the probability of wrong and correct cascades created by the immediate previous agent $n-1$. Using results from Theorem 1, the following theorem summarizes some properties of the agents' welfares.

Theorem 2. *With the same signal quality p and $\epsilon \in \mathcal{I}_k$:*

1) *The expected welfare for each agent is at least equal to the expected welfare of his predecessors. Thus, $E[\pi_n] \geq F$ and is non-decreasing in n.*
2) $\lim_{n\to\infty} E[\pi_n]$ *exists and equals:*

$$\Pi(\epsilon) = F + \frac{1}{2}\left[\frac{1}{1+\left(\frac{1-a}{a}\right)^k} - p\right]. \tag{19}$$

3) *For every agent n, $E[\pi_n]$ decreases continuously as ϵ increases over $\mathcal{I}_k$ so that:*

$$\lim_{\epsilon\downarrow\epsilon^*_k} E[\pi_n] > E[\pi_n] > \lim_{\epsilon\uparrow\epsilon^*_{k+1}} E[\pi_n] = F.$$

Proof 1) When a cascade happens, every user takes the same action and so achieves the same expected welfare. Thus, we are left to show $E[\pi_n] \geq F$ for all $n \geq k+1$. Using (18) and the form of $v^*_{0,i-1}$, $u^*_{0,i-1}$, we only need to show:

$$(1-p)a^{\frac{j+k}{2}}(1-a)^{\frac{j-k}{2}} - pa^{\frac{j-k}{2}}(1-a)^{\frac{j+k}{2}} \geq 0, \tag{20}$$

which can be seen by noting the following: since $\epsilon^*_k \leq \epsilon < \epsilon^*_{k+1}$, we have:

$$0 < \left(\frac{1-a}{a}\right)^k < \frac{1-p}{p} \leq \left(\frac{1-a}{a}\right)^{k-1} < 1. \tag{21}$$

2) For $t \in \mathbb{R}$, let $V_0(t)$ and $U_0(t)$ be the probability generating functions for the first hitting time of state k and $-k$, respectively. Using these, the limiting welfare can be written as:

$$\Pi(\epsilon) - F = \frac{1}{2}\left[(1-p)V_0(1) - pU_0(1)\right]. \tag{22}$$

Expressions for these generating functions are given in equations (47) and (48) in the Appendix D; evaluating these at $t = 1$ yields (19).

3) For a fixed p, (17) shows that F decreases in ϵ. The proof follows by using (18) and Theorem 1. □

Theorem 2 suggests a few interesting observations. First, even though a model with no observation errors yields maximum welfare for any agent, the welfare is not monotonically decreasing in ϵ. In fact, as the given ϵ approaches each ϵ_k^* from below, it is better to increase ϵ to exactly ϵ_k^* so that the welfare for each agent moves to the next local maximum. This comes from the fact that the probability of wrong cascade drops discontinuously as ϵ crosses over ϵ_k^* due to the structure of the underlying DTMC changing. Second, the first observation raises the question of when and how much observation error should be added, especially when some pre-set error is already present in this model. We will consider this question in the next section.

V. When is more noise better?

A. Asymptotic welfare

Assume that there is a fixed observation error ϵ in the model. Suppose a social planner is allowed to randomly change the history of observations with probability ϵ_s. This introduces an effect that is equivalent to increasing the observation error to

$$\epsilon_{total} = \epsilon(1-\epsilon_s) + (1-\epsilon)\epsilon_s \triangleq f(\epsilon, \epsilon_s). \tag{23}$$

Fig. 5 shows an example of the asymptotic welfare as a function of ϵ_s, when the pre-set observation error is $\epsilon = 0.15$ and $p = 0.70$. It is clear from this figure that the asymptotic welfare is maximized when $\epsilon_s \approx 0.12$, i.e., more noise is beneficial. By part 1) of Theorem 2, this also maximizes the (Cesàro) average social welfare of the entire population.

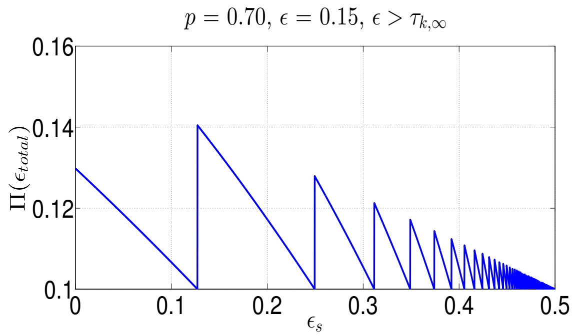

Figure 5: Asymptotic welfare as a function of additional noise.

Whether more noise is beneficial depends on the existing noise ϵ and on the signal quality p. To better understand this, for each k, define $\tau_{k,n} \in \mathcal{I}_k$ such that adding some $\epsilon_s > 0$ will improve the welfare of agent n if and only if $\epsilon > \tau_{k,n}$. Likewise, let $\tau_{k,\infty}$ denote this threshold for the asymptotic welfare. For the case in Fig. 5, since there exists a $\epsilon_s > 0$ that improves the welfare, it must be that $\epsilon > \tau_{k,\infty}$. Note also that the optimal choice of ϵ_s corresponds to a discontinuity of the welfare, which from our previous analysis corresponds to making the effective noise, ϵ_{total}, equal to ϵ^*_{k+1}. We formalize this phenomenon in the following theorem.

Theorem 3. *Assume that $\epsilon \in \mathcal{I}_k$:*

1) *The asymptotic social welfare is maximized at either $\epsilon_s = 0$ or $\epsilon_s = \frac{\epsilon^*_{k+1}-\epsilon}{1-2\epsilon} > 0$.*
2) *The latter case happens when $\epsilon^*_{k+1} > \epsilon > \tau_{k,\infty} > \epsilon^*_k$, where*

$$\tau_{k,\infty} = \frac{1}{1-2p}\left[\frac{1}{1+(\frac{1-p}{p})^{\frac{k+1}{k^2}}} - p\right]. \tag{24}$$

Proof. 1) The proof of 1) follows by part 3 of Theorem 2.
2) To prove 2), we need to find ϵ_s such that $\Pi(\epsilon) < \Pi(\epsilon_{total})$. Using (19), we obtain the lower bound $\tau_{k,\infty}$. □

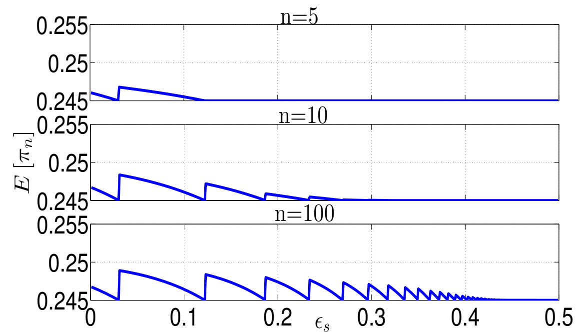

Figure 6: Expected welfare for agent n with high signal quality.

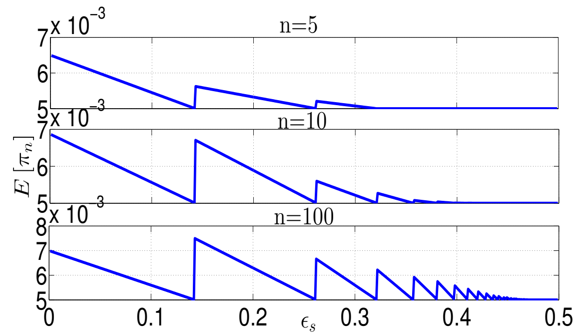

Figure 7: Expected welfare for agent n with low signal quality.

Note that when it is optimal for the social planner to add $\epsilon_s > 0$, the welfare of every individual agent is not necessarily improved. This effect is shown numerically in Fig. 6 and 7 with the same pre-set value $\epsilon = 0.15$, for high and low signal quality, respectively. These figures each shows the expected welfare for three agents corresponding to $n = 5, 10$ and 100 as a function of ϵ_s. For high signal quality ($p = 0.99$), increasing ϵ_s is beneficial for all three agents, i.e., $\epsilon > \max(\tau_{k,5}, \tau_{k,10}, \tau_{k,100})$. On the other hand, for low signal quality ($p = 0.51$), increasing ϵ_s will increase the welfare of agent $n = 100$ and decrease the welfare of agents $n = 5, 10$, i.e., $\epsilon > \tau_{k,100}$, but $\epsilon < \min(\tau_{k,5}, \tau_{k,10})$.

B. Individual welfares in the low-noise region

In Fig. 6 and 7, we observe that not all agents' welfares are improved by adding noise. In this section, we take a deeper look and consider the low-noise region, defined as when the pre-set observation error $\epsilon \in \mathcal{I}_2$. In this region, it takes only two (consecutive) agents to start a cascade

as in the noise-free model. We derive a few important results for the sequence $\{\tau_{2,n}\}_{n=1}^{\infty}$, where n is the agent index. For notational simplicity, we suppress $k = 2$ and abbreviate $\tau_{2,n}$ as τ_n. Using part 3) of Lemma 3, (18) is rewritten as:

$$E\left[\pi_n\right] = F + \frac{1}{2}\left[(1-p)\left(\frac{a}{1-a}\right)^k - p\right] u^*_{0,n-1}. \tag{25}$$

Now using part 1) and 2) of Lemma 3, if $n > k$ then $E[\pi_n] = E[\pi_{n-1}]$ is true for k and n, both odd or both even. Thus, we can simplify the welfare in (25) for the low noise region, as stated in the next lemma.

Lemma 4.

1) For $\epsilon \in \mathcal{I}_2$ and $m \geq 1$:

$$E\left[\pi_{2m+2}\right] = E\left[\pi_{2m+1}\right] \tag{26}$$

$$= F + \frac{1}{2}\left[(1-p)a^2 - p(1-a)^2\right] \frac{1 - [2a(1-a)]^m}{1 - 2a(1-a)}. \tag{27}$$

2) For $\epsilon \in \mathcal{I}_3$ and $m \geq 2$:

$$E\left[\pi_{2m+1}\right] = E\left[\pi_{2m}\right] \tag{28}$$

$$= F + \frac{1}{2}\left[(1-p)a^3 - p(1-a)^3\right] \frac{1 - [3a(1-a)]^{m-1}}{1 - 3a(1-a)}. \tag{29}$$

Proof See details in the archived version in [38]. □

We remind the reader that by Theorem 2, the expected welfare $E[\pi_n]$ is non-decreasing in n. Lemma 4 suggests that the welfare strictly increases between pairs of two consecutive agents, depending on whether k is odd or even.

Note that for $k = 2$, agents $n \leq 3$ can never benefit from adding ϵ_s. Using Lemma 4, the following theorem characterizes the sequence $\{\tau_n\}_{n=4}^{\infty}$. In particular, Theorem 4 says that the sequence $\{\tau_n\}_{n=4}^{\infty}$ can be divided into even and odd sub-sequences, where both sub-sequences are decreasing and have the same limit τ_∞ as in (24) with $k = 2$. Thus, if adding noise improves the limit welfare, then it will benefit all but the first few finite number of agents.

Theorem 4. *For $m \geq 2$:*

1) Both sub-sequences $\{\tau_{2m}\}$ *and* $\{\tau_{2m+1}\}$ *decrease and have the same limit* τ_∞*, and*

2) $\tau_{2m+1} > \tau_{2m}$.

Proof. Here, we outline the sketch of the proof for $\tau_n, n \in \{4, 5, 6, 7\}$ with a graphical illustration as in Fig. 8, which plots the welfares of these agents as a function of ϵ. To show that part 1) is true, we essentially show $\tau_4 > \tau_6$ and $\tau_5 > \tau_7$. These correspond to showing point A is higher than point C, and point B is higher than point D, respectively. To prove part 2), we essentially show that point C is higher than point D. The complete proof can be found in the archived version in [38].

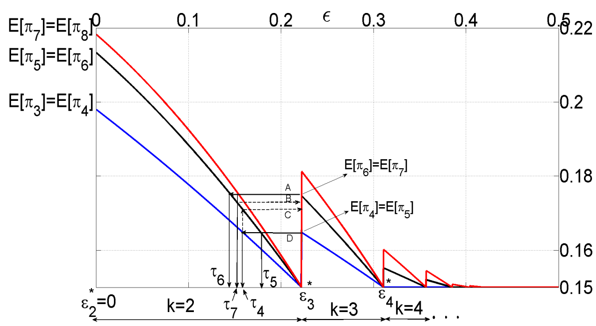

Figure 8: An illustration of ordering for $\tau_n, n \in \{4, 5, 6, 7\}$.

$\square$

Moreover, Theorem 4 also suggests an interesting fact: with the same signal quality p, $\epsilon \in \mathcal{I}_2$, and $m \geq 2$, if adding noise benefits agent $2m$ it will also benefit agents $2m + 2, 2m + 4, \ldots$; and if adding noise benefits agent $2m + 1$ then it will also benefit agents $2m$ and $2m + 3, 2m + 5, \ldots$. A consequence of this result is that while adding noise never benefits the first three agents, all other agents' welfares, including agent 4's, are better off as long as agent 5's welfare is.

VI. Adding noise versus using guinea pigs

In this section, our goal is to study the trade-offs of two schemes (mentioned below) for welfare improvements. Assume that $\epsilon \in \mathcal{I}_k$ and $\epsilon > \tau_{k,\infty}$, thus increasing the observation error to ϵ^*_{k+1} is beneficial for the average asymptotic social welfare. We refer to this as Scheme 1. Extrapolating from Theorem 4, it stands to reason that in Scheme 1 there exists a smallest agent index n_1 who benefits, and there also exists an agent index n_2 beyond which all subsequent

agents $n \geq n_2$ benefit. On the other hand, for agents $n \in (n_1, n_2)$, some benefit and some do not. The second scheme we consider is motivated by [19] where the author restricts the first $M+1$ agents (called "guinea pigs") to use only their private signals. The purpose of this scheme is to skew the initial state distribution toward a correct cascade, owing to the informativeness of the signal quality. Thus, the asymptotic social welfare is improved at the cost of the guinea pigs' payoffs. We summarize these two schemes as follows:

1) *Scheme 1*: Increase ϵ to the next threshold, ϵ^*_{k+1} (keeping $M = 1$ fixed).
2) *Scheme 2*: Sacrifice the first $M+1$ agents as guinea pigs (keeping ϵ fixed).

In Sections VI-A and VI-B below, we characterize the expected welfare for each of these schemes by using (25) and modifying the corresponding probability of wrong cascade.

A. Scheme 1: Adding observation errors

As we increase ϵ to ϵ^*_{k+1}, note that $\left(\frac{a}{1-a}\right)^k = \frac{p}{1-p}$. and $\frac{v^*_{0,n}}{u^*_{0,n}} = \left(\frac{a}{1-a}\right)^{k+1}$, for all $n > 1$. Thus, (25) becomes:

$$E[\pi_n]|_{\epsilon=\epsilon^*_{k+1}} = F + \frac{p}{2}\left[\left(\frac{p}{1-p}\right)^{1/k} - 1\right] u^*_{0,n-1}. \tag{30}$$

B. Scheme 2: Sacrificing guinea pigs

Using notation from [19], assume that there are $M+1$ guinea pigs who have to use only private signals.[11] For these agents, we have $E[\pi_n] = F$, $\forall\, n = 1, 2, \ldots, M+1$. The effect of having $M+1$ guinea pigs is equivalent to changing the distribution for the initial state of the Markov chain and then starting this chain with agent $M+2$. For each $i \in \{-M-1, -M, \ldots, M, M+1\}$, denote Δ_i as the probability that the $M+1$ guinea pigs result in $h_{M+1} = i$ for the subsequent agents $n \geq M+2$. For agents $n \geq M+2$, let $\tilde{u}^*_{i,n}$ be the probability of wrong cascade of agent n, conditioned on $h_{M+1} = i$. Their corresponding expected welfare can be written as:

$$F + \frac{1}{2}\sum_{|i|\leq M+1} \Delta_i \left[(1-p)\left(\frac{a}{1-a}\right)^k - p\right]\tilde{u}^*_{i,n-1}, \tag{31}$$

[11] In [19], the first agent always follows his own signal and so M is the number of additional agents forced to follow their own signals. In our case, more than the first agent may always follow their signals, depending on the noise, but we still use $M+1$ to denote the total number of agents that do this to be consistent with [19].

where the initial distribution for $h_{M+1} = i$ with $M + 1 - i$ even is:

$$\Delta_i = \binom{M+1}{\frac{M+1-i}{2}} a^{\frac{M+1+i}{2}} (1-a)^{\frac{M+1-i}{2}}, \tag{32}$$

and $\Delta_i = 0$ for $M + 1 - i$ odd.

For $-k < i < k$, again using techniques from [1], the probability of a wrong cascade first occurring at time $n \geq M + 2$ starting from state i is given by:

$$\tilde{u}_{i,n} = \begin{cases} 0, & \tilde{n} - (k+i) \text{ odd}, \\ \frac{2^{\tilde{n}}}{k} \left(\frac{1-a}{a}\right)^{\frac{k+i}{2}} (a(1-a))^{\frac{\tilde{n}}{2}} \tilde{B}_{k,\tilde{n}}, & \tilde{n} - (k+i) \text{ even}, \end{cases} \tag{33}$$

where $\tilde{B}_{k,\tilde{n}} = \sum_{\nu=1}^{\nu<k} \cos^{\tilde{n}-1}\left(\frac{\nu\pi}{2k}\right) \sin\left(\frac{\nu\pi}{2k}\right) \sin\left(\frac{\nu\pi(k+i)}{2k}\right)$ and $\tilde{n} = n - M - 1$. Therefore, the probability of wrong cascade for an agent n is:

$$\tilde{u}^*_{i,n} = \begin{cases} 1, & i \leq -k, \\ 0, & i \geq k, \\ 0, & \tilde{n} < k + i, \\ \sum^{\tilde{n}}_{\substack{m=k+i \\ m-(k+i)\, even}} u_{i,m}, & \tilde{n} \geq k + i. \end{cases} \tag{34}$$

Thus, revisiting (31), the summation in the welfare of every agent n can be further divided into three components corresponding to three different ranges of i: $i < -k$, $-k \leq i \leq k$, and $i > k$. Whether all these three ranges exist depends on the values of $M + 1$ and k. It is obvious that for a fixed k, letting $M + 1$ increases increases the number of terms for which $i > k$. In other words, for high enough M, the initial state after the last guinea pig is more likely equal to the correct herding state, thus improving average social welfare. However, this happens at the cost of sacrificing the first $M + 1$ agents. In the next section, we highlight this trade-off and offer an alternative solution that combines the two schemes. We refer to this as:

3) *Scheme 3:* sacrificing the first $M+1$ agents as guinea pigs, and at the same time increasing ϵ to ϵ^*_{k+1}.

C. Comparisons and trade-offs

1) Individual welfare: Figures 9-12 show numerical results for the three schemes mentioned above.

Each figure shows $E[\pi_n]$ versus n for these schemes; the choice of p and ϵ varies in different figures. In Figs. 9 and 10, the signal quality is high, and the initial noise is high (so that $k = 8$) and low ($k = 2$), respectively. Figs. 11 and 12 show analogous curves with a low signal quality. In each figure, the curves labeled $M = m$ for some integer m correspond to Scheme 2 with $m+1$ guinea pigs (where $M = 1$ is the baseline case since the first two agents always follow their own signals). Scheme 1 is shown in the curve labeled "adding noise" and Scheme 3 is shown in the curve labeled by "$M = k$ and adding noise," where again $k + 1$ is the number of guinea pigs. Note that trade-offs exist for various values of p and ϵ. For high M, Scheme 2 eventually gives better individual welfare as compared to Scheme 1. However, the larger the observation error ϵ is, the more agents that need to be sacrificed. This trade-off is alleviated by using Scheme 3, which yields better welfares while sparing almost half of the would-be guinea pigs.

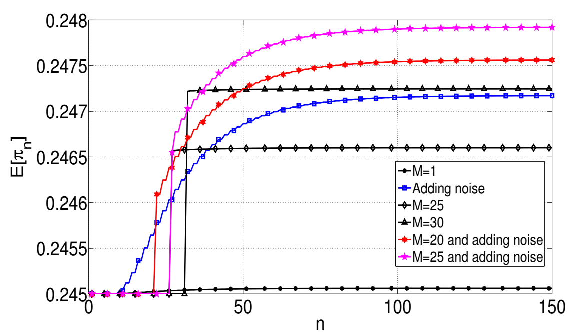

Figure 9: Agent n welfare for $p = 0.99$, $k = 8$, $\epsilon = 0.357 \in (\tau_{k,10}, \epsilon_9^*)$.

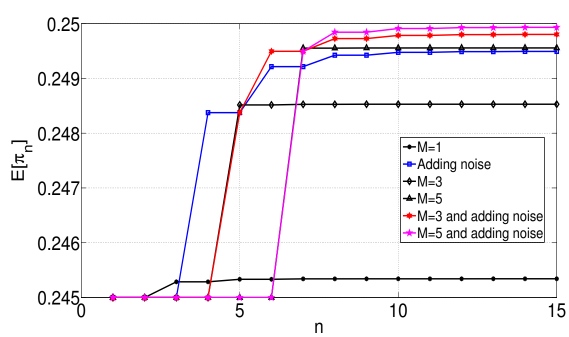

Figure 10: Agent n welfare for $p = 0.99$, $k = 2$, $\epsilon = 0.08 \in (\tau_{k,5}, \epsilon_3^*)$.

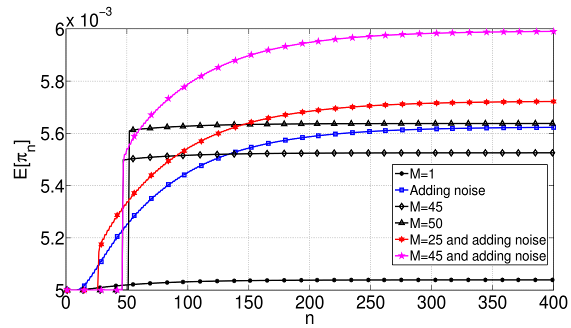

Figure 11: Agent n welfare for $p = 0.51$, $k = 8$, $\epsilon = 0.437 \in (\tau_{k,10}, \epsilon_9^*)$.

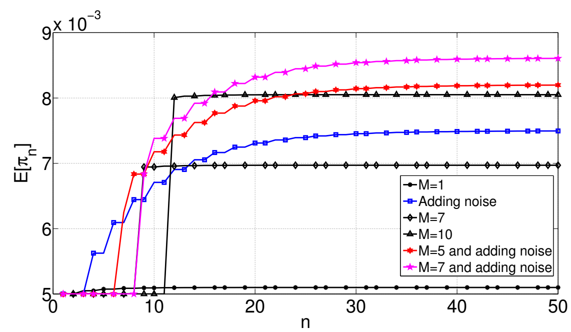

Figure 12: Agent n welfare for $p = 0.51$, $k = 2$, $\epsilon = 0.245 \in (\tau_{k,5}, \epsilon_3^*)$.

2) Asymptotic welfare: By the law of large numbers, as $M + 1$ increases, the probability that agent $M+2$ starts a correct cascade increases to 1. Thus, for large enough M, Scheme 2 achieves better asymptotic welfare as compared to Scheme 1. In this section, we aim to give a lower bound on the minimum number of guinea pigs, $M + 1$, that makes Scheme 2 asymptotically better than Scheme 1.

Proposition 1. *For $\epsilon \in \mathcal{I}_k$ and $\epsilon > \tau_{k,\infty}$ Scheme 2 achieves higher asymptotic welfare than Scheme 1 only if $M + 1 > k$.*

Proof. We refer to the case where the social planner does nothing as Scheme 0 (i.e. neither uses guinea pigs nor increases ϵ). Assume that $M + 1 \leq k$. For any $\epsilon \in \mathcal{I}_k$, by Lemma 1 all agents $n = 1, ..., k$ must use private signals even when agent n is not a guinea pig. By this argument, for a given sequence of private signals, Scheme 2 and Scheme 0 have identical samples path, thus yielding identical asymptotic welfare. Moreover, since $\epsilon > \tau_{k,\infty}$, Scheme 1 is better than

Scheme 0. Thus, the number of guinea pigs, $M + 1$ needs to be strictly higher than k to achieve better asymptotic welfare. $\square$

As a consequence of Proposition 1, as $\epsilon \to 0.5$, i.e., $a \to 0.5$, we have $k \to \infty$. Thus Scheme 2 cannot give better asymptotic welfare for a finite number of guinea pigs. For $k = 2, 3, \ldots, 10$, Figs. 13 and 14 illustrate how the required minimum M (obtained numerically) increases as ϵ approaches to 0.5 for high and low signal quality, respectively. In addition to plotting this for each given k, we also show a curve that connects the maximum optimal M for each k, which is clearly increasing in ϵ. This shows that for high noise levels, Scheme 2 is only viable if one is willing to sacrifice a large number of initial users.

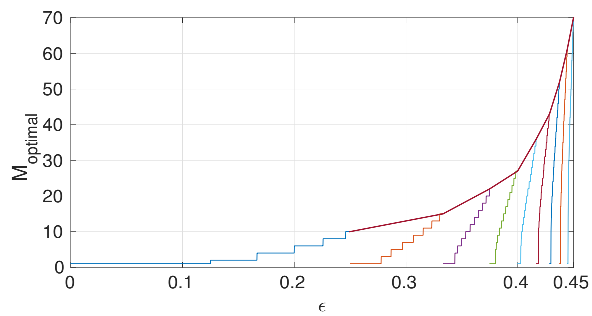

Figure 13: $p = 0.51$, from left to right: $k = 2, 3, ..., 10$

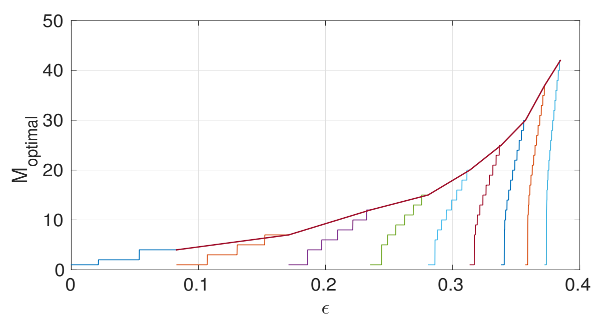

Figure 14: $p = 0.99$, from left to right: $k = 2, 3, ..., 10$

VII. Cascades with no common database

So far, our model has assumed that there is observation noise by letting the agents report their actions on a common database which all other agents can see. Even though the reporting process is in error with probability ϵ, the observation errors are common knowledge. In other

words, what one agent observes reflects exactly what any of the previous agents observed. In some settings, this may be a wrong assumption, e.g., if agents may pull out information using different methods or through different channels. Additionally, such a common database may not be valid in other social networks settings different from online recommendation systems. In this section, we step away from this assumption and look at in some sense the opposite case. In particular, we consider the scenario when each agent individually observes an erroneous realization of the previous agents' true actions, where all errors are assumed to be *i.i.d.* Again, the distribution of the errors is still common knowledge. In such a scenario, not only are the agents blind to other agents' private signals, but they also do not know what their predecessors have observed.

A. Will cascades persist?

We remind the reader that the key to previous analysis of cascades with a common database is that every subsequent agent knows the exact answers to two questions: 1) Has a cascade started yet? and 2) If the answer to the previous question is Yes, then when did that cascade start? The reasons are two-fold: 1) The distributions of the prior knowledge about V, the private signal, and the observation errors are given; and 2) A common observation database is available to all agents. Without the second assumption, no agent can infer if and when herding has started.

Again, since the pre-cascading actions are independent, Property 1 still holds as in the common database scenario. However, the following property shows that, without a common observation database, herding is not recurrent.

Property 2′ *Without the common database, a cascade needs not be permanent.*

Proof Assume that the statement is not true. We will provide a counter-example to this using just four agents. Denote O^i_j as the observations of agent i about the action A_j. Again, assume that O^i_j is a noisy version of A_j, so that O^i_j differs from A_j in an *i.i.d.* manner with error probability ϵ for all agents i and their predecessors $1 \leq j \leq i-1$. Now the information set of each agent i includes his private signal S_i and his observations $O^i_1, \ldots, O^i_{i-1}$ of the true actions $A_1, \ldots, A_{i-1}$.

The decision rule for agents 1 and 2 is to follow their private information, the same as before. Since the first two agents both use their own signals, agent 3 takes account of the independence of O^3_1, O^3_2. Thus, agent 3's decision rule remains the same: herds only if O^3_1, O^3_2 are identical and $\epsilon < \epsilon^*_3$. Now assuming that $\epsilon < \epsilon^*_3$, the Bayesian update of agent 4 is:

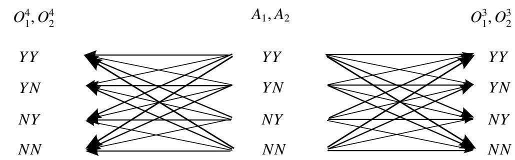

Figure 15: Different realizations of observations history for agents 3 and 4.

$$\gamma_4 = \frac{\mathbb{P}[O_1^4, O_2^4, O_3^4 | V = G]\mathbb{P}[S_4 | V = G]}{\text{numerator} + \mathbb{P}[O_1^4, O_2^4, O_3^4 | V = B]\mathbb{P}[S_4 | V = B]}. \tag{35}$$

Now, conditioning on $V \in \{G, B\}$, we can write:

$$\begin{aligned}\mathbb{P}[O_1^4, O_2^4, O_3^4] &= \mathbb{P}[O_3^4 | O_1^4, O_2^4]\mathbb{P}[O_1^4, O_2^4] \\ &= (1 - \epsilon)\mathbb{P}[A_3 = Y, O_1^4, O_2^4] + \epsilon\mathbb{P}[A_3 = N, O_1^4, O_2^4], \end{aligned} \tag{36}$$

where (36) holds if $O_3^4 = Y$. For $O_3^4 = N$, ϵ and $1 - \epsilon$ are interchanged. The calculation of γ_4 is possible if we can evaluate $\mathbb{P}[A_3, O_1^4, O_2^4]$ for both $A_3 = Y, N$ and $V = G, B$. This can be done by conditioning on the 4 possibles realizations of agent 3's observations as illustrated in Fig. 15.

As an example, consider the case the information set of agent 4 is $\left\{S_4, O_1^4, O_2^4, O_3^4\right\} = \{0, Y, Y, N\}$. The above procedure gives:

$$\gamma_4 = \frac{\left[\epsilon\alpha_G + (1 - \epsilon)\beta_G + 2(1 - a)\delta_G\right](1 - p)}{\text{numerator} + \left[\epsilon\alpha_B + (1 - \epsilon)\beta_B + 2a\delta_B\right] p}, \tag{37}$$

where subscripts correspond to true value $V \in \{G, B\}$, and: $\alpha_G = [a - \epsilon(1 - \epsilon)]^2$, $\alpha_B = [(1 - a) - \epsilon(1 - \epsilon)]^2$, $\beta_G = \beta_B = [\epsilon(1 - \epsilon)]^2$, $\delta_G = \epsilon(1 - \epsilon)[a - \epsilon(1 - \epsilon)]$, $\delta_B = \epsilon(1 - \epsilon)[(1 - a) - \epsilon(1 - \epsilon)]$. Based on the decision rule in (2), one would need to compare γ_4 and $1/2$. Using the assumption that $0 \leq \epsilon < \epsilon_3^*$, which means $\frac{a}{1-a} \leq \frac{p}{1-p} < \left(\frac{a}{1-a}\right)^2$, we deduce $\gamma_4 < 1/2$, thus agent 4 follows his own signal and chooses N. Applying similar procedure for the remaining realizations of agent 4's observation history, we have the decision rule shown in Table I.

Thus agent 4 herds only when all of his three observations are identical. Now given that agent 3 cascades to the first two agents' action, agent 4 would choose the same action only if there were no observation errors, i.e., with probability $(1 - \epsilon)^3$, which is bounded away from 1. Thus in this scenario the cascade does not persist with probability one. □

$\{S_4, O_1^4, O_2^4, O_3^4\}$	A_4	$\{S_4, O_1^4, O_2^4, O_3^4\}$	A_4
$1YYY$	Y	$1YYN$	Y
$0YYY$	Y	$0YYN$	N
$1YNY$	Y	$1YNN$	Y
$0YNY$	N	$0YNN$	N
$1NYY$	Y	$1NYN$	Y
$0NYY$	N	$0NYN$	N
$1NNY$	Y	$1NNN$	N
$0NNY$	N	$0NNN$	N

Table I: Decision rule of agent 4.

B. *Complexity analysis*

We now discuss the complexity of updating the posterior probability for an arbitrary agent n without a common database. Assume that $\epsilon \in \mathcal{I}_k$ as before. We have the following:
1) For agents $n = 1, \ldots, k$, $|h_{n-1}| < k$ thus cascades cannot happen. The decision rule is simple: always follow one's signal. Thus complexity is $O(1)$.
2) However, any agent $n > k$ faces the possibility that some prior agent was in a cascade. While agent n's decision rule remains the same as in (2), the complexity of his Bayesian update of the posterior probability, γ_n, increases exponentially as n grows. In fact, conditioned on his own observations sequence, every agent n needs to account for 2^{n-2} possible observations sequences of agent $n-1$. For each one of those 2^{n-2} possibilities, agent n needs to know the decision rule of agent $n-1$, which depends on another 2^{n-3} possible realizations of agent $n-2$'s history. Since the pre-herding agent $n = k$'s update is of order $O(k)$ (including first calculating the sufficient statistics h_{n-1} and then comparing that to $\pm k$), the complexity of the update for agent $n > k$ is:

$$\left(\prod_{m=k}^{n-1} 2^m\right) O(k) = O(k2^{n^2-k^2}), \tag{38}$$

which blows up exponentially as $(n-k)$ increases.

For comparison, we also evaluate the complexity when a common database is available. For agents $n = 1, .., k$, the complexity is also $O(1)$ by simply checking if $n \leq k$. For any agent $n > k$, the decision process is done by simply computing the sufficient history statistic h_{n-1} and comparing that with $\pm k$ at each step. Note that $|h_{n-1}|$ never exceeds k throughout the updating

process, and so in the worst case agent n endures a complexity of $O(n)$. This illustrates a benefit of having a common online platform for aggregating information. On one hand, such a platform lets agents get more information than they probably could by just having their own observations. On the other hand, it also provides a computational advantage such that if there is noise present, having this common platform makes it easier for agents to learn. The complexity incurred without a common database for Bayes-rational agents suggests that in such case, it is worth considering non-Bayesian approaches in the same spirit as [26]. We leave such approaches for future work.

VIII. Other Generalizations

In this section, we discuss several other generalizations of our model. We start next with examining different tie-breaking rules.

A. Comparison of different tie-breaking rules

When an agent's postierior probability satisfies $\gamma_n = 1/2$, its expected pay-off under either action is the same. In [9], the authors assumed that when such a tie occurs, agents randomize their choices. Another option for breaking a tie is to follow the previous action whenever a tie happens. In our model, we instead assume that an agent chooses to follow his own private signal when he is indifferent. Our assumption not only simplifies the analysis, but as we state in the next proposition, also is the better tie-breaking rule as compared to the other two: by following their own signals, agents who face a tie have lower wrong cascade probabilities and therefore, higher welfare.

Proposition 2. *For a given integer $k \geq 2$:*

1) *When $\epsilon \in (\epsilon_k^*, \epsilon_{k+1}^*)$, all agents are never in a tie. Thus all three tie-breaking rules perform equivalently.*
2) *When $\epsilon = \epsilon_k^*$, agents who face a tie achieve the lowest (resp. highest) wrong cascade probability if they follow their own signals (resp. follow the previous action).*

Proof. We start by reminding the reader that a tie happens when the posterior probability satisfies $\gamma_n = 1/2$, i.e., there exists an integer $k \geq 2$ such that $\left(\frac{a}{1-a}\right)^k = \frac{p}{1-p}$. To prove the first part, note that when $\epsilon \in (\epsilon_k^*, \epsilon_{k+1}^*)$, the above equality cannot happen. Thus all tie-breaking rules give the same outcome.

To prove the second part, assume that $\epsilon = \epsilon_k^*$ for some $k \geq 2$. If an agent n encounters a history with $|h_{n-1}| = k - 1$ that is opposite to his private signal, he has to resort to the tie-breaking condition. If he follows the previous action, a cascade happens where the absorbing states of the underlying Markov chain are $\pm(k - 1)$. On the other hand, if he follows his private signal, a cascade does not happen yet, as the absorbing states of the underlying Markov chain are $\pm k$. Using the equations (12)-(8), the latter cases give a lower wrong cascade probability and higher correct cascade probability. As a consequence by (18), when agent n follows his own signal, he has a higher welfare. Moreover, this also yields a higher asymptotic welfare as seen from (19). This completes the proof. □

The intuition behind this result is as follows: when the tie-breaking rule is to follow an agent's private signal, the underlying Markov chain has the two absorbing states expanded by one unit on both sides as compared to when following the previous action. This induces a longer expected time until a cascade event (on either side). However, in both scenarios the chains have the same drift, which is always in favor of the correct cascade. This effect decreases (increases) the wrong (correct) cascade probabilities. This can also be seen in Fig. 4: each point of discontinuity happens exactly when k is increased by 1 unit where the probability of wrong cascades drops. If agents follow the previous actions, this yields the upper end of the discontinuity; while if agents follow their private signals, this yields the lower end of the discontinuity. Moreover, any randomized tie-breaking rule results in a linear combination of these two end points, which also yields a higher wrong cascade probability than when agents follow their own signals.

B. Cascades with bounded rationality

Next, we consider the possibility that agents might make mistakes in choosing their optimal action. This captures a more realistic assumption since it is well documented that real-world agents are not always rational. In fact, the reasons underlying human irrationality have been well argued in [20] where the authors presented the numerous cognitive deficiencies that could account for human systematic deviation from the perceived normative behaviors. Further, from the viewpoint of Prospect Theory ([5]), the action error can also be argued to be a consequence of inconsistency in preferences when agents tend to act differently under the possibilities of gains (facing a correct herding) and losses (being in a wrong herd). Finally, such a deviation from full rationality has also been used for equilibrium selection, in particular, for trembling-hand

perfection equilibrium ([4]).

Indeed, there is a long history of using noise to model imperfect rationality (see e.g. [8], [14]) resulting in agents having so-called "noisy best responses." Here, we follow the approach in those works, and examine a simple form of bounded rationality. In particular, assume that all agents derive the exact posterior probability of the true value of the item, but then choose the sub-optimal action with probability $\epsilon_a \in (0, 0.5)$; we call this an "action error." Assume that this probability is known to all other agents, and is the same across all agents. Again, as before, assume that the observation database is still common knowledge, with *i.i.d.* observation errors with probability $\epsilon \in (0, 0.5)$.

The analysis is similar as in the setting with no action errors, so we omit the details. The reason being we can combine the two types of errors, ϵ and ϵ_a, to an equivalent error type, $f(\epsilon, \epsilon_a)$, and use that in updating the posterior probability as before. As a result, the probabilities of wrong and correct cascades, $u^*_{0,n}$ and $v^*_{0,n}$, respectively, are as those in a model with only observation errors of rate $f(\epsilon, \epsilon_a)$. However, the agent welfares are now reduced due to the probability of choosing the non-optimal action, ϵ_a. In particular, for agents $1 \leq n \leq k$, they all have the same welfare given by:

$$\begin{aligned} E[\pi_n] &= \{\mathbb{P}[A_n = Y | V = G] - \mathbb{P}[A_n = Y | V = B]\}/4 \\ &= (1 - 2\epsilon_a)(2p - 1)/4 = (2a_1 - 1)/4 \triangleq F_1. \end{aligned} \tag{39}$$

where $a_1 = f(\epsilon_a, p)$.

For agents $n \geq k + 1$:

$$\begin{aligned} E[\pi_n] &= \{\mathbb{P}[A_n = Y | V = G] - \mathbb{P}[A_n = Y | V = B]\}/4 \\ &= F_1 + \left[(1 - a_1 - \epsilon_a)v^*_{0,n-1} - (a_1 - \epsilon_a)u^*_{0,n-1}\right]/2 \\ &= F_1 + (1 - 2\epsilon_a)\left[\frac{1-p}{2}v^*_{0,n-1} - \frac{p}{2}u^*_{0,n-1}\right]. \end{aligned} \tag{40}$$

By comparing (39) and (40) against (17) and (18), it is clear that the action error ϵ_a only reduces the welfare of every agent by a factor of $1 - 2\epsilon_a$ as compared to a model with only observation errors with probability $f(\epsilon, \epsilon_a)$. Thus, the expected welfare of agents in a quasi-rational model inherits many properties from the fully rational setting with noisy observations as in Theorem 2 (e.g., the overall welfare is not monotonic in the error rate).

C. General prior for the true value, V

In this section, we consider the scenario where the binary true value V of the item might not be equally likely, but the *ex-ante* pay-off is still restricted to be zero.[12] In particular, assume that $\mathbb{P}[V = G] = \lambda \in (0, 1)$. Let the rewards (resp. losses) of buying when $V = G$ (resp. $V = B$) be x (resp. $-y$) where $x, y \geq 0$ and $x + y > 0$. In all previous sections, we have assumed $x = 1, y = 0$ and $\lambda = 1/2$. In this general scenario, setting the *ex-ante* payoff to 0 leads to the new cost of the item being $C = \lambda x - (1 - \lambda)y$. Putting these all together, the general payoff is given by:

$$\pi_n = \begin{cases} 0, \text{ if } A_n = N, \\ x - C = (1 - \lambda)(x + y) \text{ if } A_n = Y, V = G, \text{ and} \\ -y - C = -\lambda(x + y) \text{ if } A_n = Y, V = B. \end{cases} \tag{41}$$

The ex-post expected payoff of agent n when taking an action $A_n \in \{Y, N\}$ is calculated as:

$$\begin{cases} E\left[\pi_n | A_n = N\right] = 0, \text{ and} \\ E\left[\pi_n | A_n = Y\right] = (x + y)(\gamma_n - \lambda). \end{cases} \tag{42}$$

By (42), the new decision rule is based on comparing the posterior probability, γ_n, to λ. However, as λ changes, the posterior probability is a function of λ and is given by:

$$\gamma_n = \mathbb{P}[V = G | S_n, \mathcal{H}_n] = \frac{1}{1 + \beta_n \ell_{n-1} \frac{1-\lambda}{\lambda}} \tag{43}$$

Thus, the new (user) optimal decision rule compares $\beta_n \ell_{n-1}$ and 1. Moreover, since β_n, ℓ_{n-1} are conditioned on V and, thus, do not change as λ changes, the new decision rule with respect to $\beta_n \ell_{n-1}$ is the same as before. Conditioned on V, the transition probabilities of the underlying Markov chain are identical to the case $\lambda = 1/2$. As a result, conditioned on V, the probabilities of cascades for an agent n are also given by:

$$\begin{cases} \mathbb{P}\left[wrong | V = G\right] = u^*_{0,n} = 1 - v^*_{0,n}, \text{ and} \\ \mathbb{P}\left[wrong | V = B\right] = v^*_{0,n} = 1 - u^*_{0,n}. \end{cases} \tag{44}$$

However, the (unconditional) probabilities of wrong and correct cascades are given as $\lambda u^*_{0,n} + (1 - \lambda)v^*_{0,n}$ and $\lambda v^*_{0,n} + (1 - \lambda)u^*_{0,n}$, respectively. Given that λ, x, y are fixed, the welfare for agents

[12]The case of zero *ex-ante* payoff is the interesting case as at the outset agents have no preference for either action, and thus the additional information from the signals and past observations could be helpful.

are scaled by a constant factor as compared to when V is equally likely G or B. We show this next. All agents $1 \leq n \leq k$, who use their own signals, have the same payoff:

$$
\begin{aligned}
E[\pi_n] &= (1-\lambda)E[\pi_n|V=B] + \lambda E[\pi_n|V=G] \\
&= (1-\lambda)(-\lambda)(x+y))P[A_n = Y|V=B] \\
&+ \lambda(1-\lambda)(x+y)P[A_n = Y|V=G] \\
&= \lambda(1-\lambda)(x+y)(2p-1) \stackrel{\triangle}{=} F_2. \qquad (45)
\end{aligned}
$$

For agents $n \geq k+1$, who face the possibility of cascades:

$$
\begin{aligned}
&E[\pi_n] = \lambda(1-\lambda)(x+y)\{\mathbb{P}[A_n = Y|V=G] - \mathbb{P}[A_n = Y|V=B]\} \\
&= \lambda(1-\lambda)(x+y)\left[v^*_{0,n-1} + p(1 - v^*_{0,n-1} - u^*_{0,n-1})\right] \\
&- \lambda(1-\lambda)(x+y)\left[u^*_{0,n-1} + (1-p)(1 - u^*_{0,n-1} - v^*_{0,n-1})\right] \\
&= F_2 + 2\lambda(1-\lambda)(x+y)\left[v^*_{0,n-1} - u^*_{0,n-1}\right]. \qquad (46)
\end{aligned}
$$

Thus, comparing (45) and (46) to (17) and (18), the welfare of all agents are scaled by a factor of $4\lambda(1-\lambda)(x+y)$. If $x+y>0$ is fixed while λ can change in $(0,1)$, the welfare are maximized if the true value V is equally likely G or B, i.e. $\lambda = 1/2$.

IX. Conclusions and future work

This paper studied the effect of observation error in a simple Bayesian information cascade where agents can be either rational or quasi-rational. We showed that despite the presence of observation errors, cascades happen in finite time with probability 1. However, observation errors can delay the onset of a cascade; we also determined the error thresholds that increase this delay. In addition, cascades with observation errors have the same level of fragility as when such errors are not present, and in both cases they can be broken by an additional private signal. Using a Markov chain based analysis we determined the probabilities of cascades for an arbitrary agent and used these to calculate the agents' welfare based on the given signal quality and the error. Our main result shows that for certain ranges of parameters, adding a controlled amount of observation error can lead to higher welfare for all but a finite number of agents. In such scenarios, we compared and contrasted the trade-offs of different methods for improving the agents' welfares. These results may be helpful for a platform operator who has the option to

weigh the trade-offs and choose the best method that would benefit his platform. Our results are strongly based on the existence a common database, which stores all the erroneous observations of past actions. Moreover, we argued that such a database provides an important computational benefit to agents in performing Bayesian updates. In future work we plan on generalizing to agents with heterogeneous private signals and observation errors, and allowing each agent to only observe subsets of past agents' actions prior to taking actions.

References

[1] W. Feller, *An introduction to Probability Theory and Its Applications*, vol. I, 3rd ed., New York Wiley, 1968.

[2] T. Cover, *Hypothesis Testing with finite statistics*, The Annals of Mathematical Statistics, vol. 40, No. 3, pp. 828-835, 1969.

[3] M. Hellman, *Learning with finite memory*, Ph.D. Dissertation, Stanford University, 1969.

[4] R. Selten, *A reexamination of the perfectness concept for equilibrium points in extensive games*, International Journal of Game Theory, vol. 4, pp. 25-55, 1975.

[5] D. Kahneman, A. Tversky, *Prospect Theory: An Analysis of Decision under Risk*, Econometrica, vol. 47, issue 2, pp. 263-292, 1979.

[6] D. Stoyan, *Comparison Methods for Queues and Other Stochastic Models*, John Wiley and Sons, 1983.

[7] T.L. Lai, H. Robbins, *Asymptotically efficient adaptive allocation rules*, Advances in applied mathematics, vol. 6, pp. 4-22, 1985.

[8] R. Rosenthal, *A Bounded Rationality Approach to the Study of Noncooperative Games*, International Journal of Game Theory, vol. 18, pp. 273-92, 1989.

[9] S. Bikhchandani, D. Hirshleifer, and I. Welch. *A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades*, J. Polit. Econ., vol. 100, No. 5, pp. 992-1026, 1992.

[10] A. Banerjee, *A Simple Model of Herd Behavior*, The Quarterly Journal of Economics, vol. 107, pp. 797-817, 1992.

[11] I. Welch, *Sequential Sales, Learning, and Cascades*, Journal of Finance, vol. 47, issue 2, pp. 695-732, 1992.

[12] I. Lee, *On the Convergence of Informational Cascades*, Journal of Economic Theory, vol. 61, pp. 396-411, 1993.

[13] R. Agrawal, *Sample mean based index policies with* $O(logn)$ *regret for the multi-armed bandit problem*. Advances in Applied mathematics, vol. 27, pp. 1054-1078, 1995.

[14] R. McKelvey, T. Palfrey, *Quantal Response Equilibria for Normal Form Games*, Games and Econ Behavior, vol. 10, pp. 6-38, 1995.

[15] S. Bikhchandani, D. Hirshleifer, and I. Welch. *Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades*, Journal of Economic Perspectives, vol. 12, No. 3, pp. 151-170, 1998.

[16] V. Bala, S. Goyal, *Learning from Neighbours*, Review of Economic Studies, Wiley Blackwell, vol. 65, issue 3, pp. 595-621, 1998.

[17] L. Smith, P. Sorensen, *Pathological Outcomes of Observational Learning*, Econometrica, vol. 68, pp. 371-398, 2000.

[18] P. Sorensen, *Rational Social Learning*, Ph.D. Dissertation, Massachusetts Institute of Technology, 1996.

[19] D. Sgroi, *Optimizing Information in the Herd: Guinea Pigs, Profits, and Welfare*, Games and Economic Behaviour, vol. 39, pp. 137-166, 2002.

[20] E. Shafir, L. Robyn, *Rationality*, American Review of Psychology, vol 53, pp 491-517, 2002.

[21] F. Dalmagro, J. Jimenez, R. Jimenez and H. Lugo, *Co-evolution of bounded rational agents in adaptive social networks*, No. 354, Computing in Economics and Finance, 2005.

[22] J. Surowiecki, *The Wisdom of Crowds*, New York: Anchor Books, 2005.

[23] D. Ariely, *Predictably Irrational: The Hidden Forces That Shape Our Decisions*, Revised and Expanded Edition, Harper Perennial, 2010.

[24] S. Veeraraghavan, L. Debo, *Herding in queues with waiting costs: Rationality and regret*, Manufacturing & Service Operations Management, vol. 13, issue 3, pp. 329-346, 2011.

[25] D. Acemoglu, M. Dahleh, I. Lobel, and A. Ozdaglar, *Bayesian learning in social networks*, Review of Economic Studies, vol. 78, pp. 1201-1236, 2011.

[26] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, *Non-Bayesian social learning*, Games and Economic Behavior, vol. 76, pp. 210-225, 2012.

[27] D. McFadden, *The new science of pleasure: consumer choice behavior and the measurement of well-being*, NBER, Working paper No. 18687, 2013.

[28] R. Greenwood, A. Shleifer, *Expectations of Returns and Expected Returns*, NBER Working Paper No. 18686, 2013.

[29] Z. Zhang, E. Chong, A. Pezeshki, W. Moran, *Hypothesis Testing in Feedforward Networks With Broadcast Failures*, JSTSP, vol. 7, issue 5, pp. 797-810, 2013.

[30] K. Drakopoulos, A. Ozdaglar, J. N. Tsitsiklis, *On Learning With Finite Memory*, IEEE Transactions on Information Theory, vol. 59, issue 10, pp. 6859-6872, 2013.

[31] A. Gionis, E. Terzi, P. Tsaparas, *Opinion maximization in social networks*, SIAM International Conference on Data Mining (SDM), pp. 387-395, 2013.

[32] M. Förster, A. Mauleon, V. Vannetelbosch, *Trust and manipulation in social networks*, Network Science, vol. 4, issue 2, pp. 216-243, 2014.

[33] T. Le, V. Subramanian, R. Berry, *The Value of Noise for Informational Cascades*, IEEE ISIT Proceeding, pp. 1101-1105, 2014.

[34] T. Le, V. Subramanian, R. Berry, *The Impact of Observation and Action Errors on Informational Cascades*, IEEE CDC Proceeding, pp. 1917-1922, 2014.

[35] Y. Wang, P. Djuric, *Social learning with Bayesian agents and random decision making*, IEEE Transac. On Signal Processing, vol. 63, no. 12, 2015.

[36] Y. Song, *Social Learning with Endogenous Network Formation*, submitted to Journal of Economic Theory, 2015.

[37] S. Etesami, S. Bolouki, A. Nedić, T. Başar, *Conformity versus Manipulation in Reputation Systems*, IEEE CDC proceeding, pp. 4451-4456, 2016.

[38] T. Le, V. Subramanian, R. Berry, *Information Cascades with Noise*, longer archived version of this paper, available at: https://archive.org/details/info_cascades_noise_journal

Appendix A
Proof of Property 1

We show the following equivalent claim: if a cascade has not happened yet, then opposite observations cancel out. Consider an arbitrary, non-cascading agent n. Since any agent $i \in \{1, 2, ..., n-1\}$ uses his own private signal for the decision, we have: $\mathbb{P}[O_i = Y|V = 1] = \mathbb{P}[O_i =$

$N|V=0]=a$, $\mathbb{P}[O_i=N|V=1]=\mathbb{P}[O_i=Y|V=0]=1-a$. Since $O_1, O_2, ..., O_{n-1}$ are mutually independent, and independent of S_n, the posterior probability of agent n is:

$$\gamma_n(S_n, O_1, ..., O_{n-1}) = \mathbb{P}[V=G|S_n, O_1, ..., O_{n-1}]$$
$$= \frac{\prod_{i=1}^{n-1}\mathbb{P}[O_i|V=G]\mathbb{P}[S_n|V=G]}{\text{numerator} + \prod_{i=1}^{n-1}\mathbb{P}[O_i|V=B]\mathbb{P}[S_n|V=B]}$$

where all the opposite observations cancel out in both numerator and denominator. Similar reasoning holds for the case $V=B$. □

Appendix B
Proof of Property 2

The key is that since the observation database is available to all agents, every agent knows whether and, if yes, when a cascade happens. Thus, the observations from the onset of a cascade provide no information to subsequent agents, and are ignored. Moreover, since all private signals and noisy observations are *i.i.d.*, subsequent agents have identical posterior as the agent who started the cascade. Now since all agents have the same preferences, they also have the same optimal action of cascading. Thus the cascade lasts forever. □

Appendix C
Proof of Lemma 1

It is sufficient to derive ϵ_k^* from the scenario $n=k$ and $\{S_n, \mathcal{H}_n\} = \{0, Y, Y, ..., Y\}$ and we will do this by induction. For agent $n=2$ with the information set $\{0, Y\}$, Lemma 1 claims that agent 2 does not herd as long as $\epsilon \geq 0$. Indeed:

$$\gamma_2(0, Y) = \frac{(1-p)a}{(1-p)a + p(1-a)} \leq 1/2 \ \forall \epsilon \geq 0. \tag{47}$$

By the decision rule in (2), agent 2 chooses N and does not start a cascade. For agent $n=k>2$ with information set $\{0, Y, Y, ..., Y\}$, assume that $\epsilon \geq \epsilon_{k-1}^*$, i.e., all previous $k-1$ agents do not cascade. Similarly, using Bayes' rule gives:

$$\gamma_n(0, Y, ..., Y) = \frac{(1-p)a^{k-1}}{(1-p)a^{k-1} + p(1-a)^{k-1}}. \tag{48}$$

Agent n does not cascade if $\gamma_n(0, Y, ..., Y) \leq 1/2$, i.e., $(1-p)a^{k-1} \leq p(1-a)^{k-1}$. Now, since $\frac{1-a}{a} = \frac{(1-\epsilon)(1-p)+\epsilon p}{(1-\epsilon)p+\epsilon(1-p)} = \frac{(1-\epsilon)\alpha+\epsilon}{(1-\epsilon)+\epsilon\alpha}$, we arrive at the form of (5) as in Lemma 1. □

Appendix D
Proof of Lemma 2

The proof follows using techniques from [1]. Let $T_{-k,i}$ and $T_{k,i}$ be random variables denoting the first time the Markov chain hits the absorbing states $-k$ and k, respectively, starting from state i. Thus, $u_{i,n} = P\left[T_{-k,i} = n\right]$, $v_{i,n} = P\left[T_{k,i} = n\right]$. Let $U_i(t), V_i(t)$ be the corresponding probability generating functions, where $t \in \mathbb{R}$. We have: $U_i(t) = E[t^{T_{-k,i}}] = \sum_{n=0}^{\infty} u_{i,n}t^n$, $V_i(t) = E[t^{T_{k,i}}] = \sum_{n=0}^{\infty} v_{i,n}t^n$. With one-step transition probabilities being a and $1-a$, we have the the difference equations:

$$U_i(t) = atU_{i+1}(t) + (1-a)tU_{i-1}(t), \tag{49}$$

$$V_i(t) = (1-a)tV_{i+1}(t) + atV_{i-1}(t), \tag{50}$$

where $-k < i < k$, with the boundary conditions: $U_{-k}(t) = 1, U_k(t) = 0, V_{-k}(t) = 0, V_k(t) = 1$. The solutions are:

$$U_i(t) = \left[\lambda_1^{i+k}(t)\lambda_2^{2k}(t) - \lambda_1^{2k}(t)\lambda_2^{i+k}(t)\right] / \left[\lambda_2^{2k}(t) - \lambda_1^{2k}(t)\right], \tag{51}$$

$$V_i(t) = \left[\lambda_1^{i+k}(t) - \lambda_2^{i+k}(t)\right] / \left[\lambda_1^{2k}(t) - \lambda_2^{2k}(t)\right], \tag{52}$$

where $\lambda_{1,2}(t) = \left[1 \pm \sqrt{1-4a(1-a)t^2}\right] / (2at)$. Since the Markov chain starts at state $i = 0$, we find $u_{0,n}$ and $v_{0,n}$ by:

$$u_{0,n} = \frac{\mathrm{d}^n U_0(t)}{n!(\mathrm{d}t)^n}|_{t=0},\ v_{0,n} = \frac{\mathrm{d}^n V_0(t)}{n!(\mathrm{d}t)^n}|_{t=0}, \tag{53}$$

which can be written in closed-form as in (10) and (11). □

Appendix E
Coupling method for the proof of Theorem 1

Here we provide more details on the coupling method used to show that the inequalities of cascade probabilities are strict in the proof of Theorem 1. Using Proposition 1.10.4 in [6], there exist two random variables Z^1_{n-1} and Z^2_{n-1} on a common probability space such that $\mathbb{P}(Z'_{n-1} = k) = \mathbb{P}(Z^1_{n-1} = k)$, $\mathbb{P}(Z''_{n-1} = k) = \mathbb{P}(Z^2_{n-1} = k)$ and the stochastic order also holds for Z^1_{n-1} and Z^2_{n-1}

almost surely, i.e., $Z^1_{n-1} \geq Z^2_{n-1}$ *a.s*. Define two independent random variables A^1_n and B_n that are also independent of Z^1_{n-1} and Z^2_{n-1} by:

$$A^1_n = \begin{cases} 1, & \text{with probability } a^{'}, \\ -1, & \text{with probability } 1 - a^{'}. \end{cases} \tag{54}$$

$$B_n = \begin{cases} 0, & \text{with probability } \frac{a^{''}}{a^{'}}, \\ -2, & \text{with probability } 1 - \frac{a^{''}}{a^{'}}. \end{cases} \tag{55}$$

Next, define the random variable A^2_n by:

$$A^2_n = A^1_n + \mathbb{1}_{A^1_n = 1} B_n. \tag{56}$$

By (54)-(56), $A^1_n \geq A^2_n$ and A^2_n has the following distribution:

$$A^2_n = \begin{cases} 1, & \text{with probability } a^{''}, \\ -1, & \text{with probability } 1 - a^{''}. \end{cases} \tag{57}$$

Using the above, define Z^j_n for $j = 1, 2$ as:

$$Z^j_n = \begin{cases} Z^j_{n-1}, & \text{if } |Z^j_{n-1}| = k, \\ Z^j_{n-1} + A^j_{n-1}, & \text{otherwise.} \end{cases} \tag{58}$$

Thus, Z^1_n and Z^2_n have the same distribution as $Z^{'}_n$ and $Z^{''}_n$, respectively. Now, $\mathbb{P}r(Z^2_n = k)$ can be written as:

$$\mathbb{P}(Z^2_{n-1} = k) + \mathbb{P}(Z^2_{n-1} = k - 1, A^2_n = 1). \tag{59}$$

Since $Z^1_{n-1} \geq Z^2_{n-1}$, $Z^2_{n-1} = k$ implies $Z^1_{n-1} = k$. Moreover, $Z^2_{n-1} = k - 1$ also implies $Z^1_{n-1} = k - 1$ or $Z^1_{n-1} = k$. Thus (59) can be decomposed as:

$$\begin{aligned} &\mathbb{P}(Z^2_{n-1} = k, Z^1_{n-1} = k) + \mathbb{P}(Z^2_{n-1} = k - 1, Z^1_{n-1} = k, A^2_n = 1) \\ &+ \mathbb{P}(Z^2_{n-1} = k - 1, Z^1_{n-1} = k - 1, A^2_n = 1). \end{aligned} \tag{60}$$

Since $A^2_n = 1$ implies $A^1_n = 1$, (60) is smaller than:

$$\begin{aligned} &\mathbb{P}(Z^2_{n-1} = k, Z^1_{n-1} = k) + \mathbb{P}(Z^2_{n-1} = k - 1, Z^1_{n-1} = k, A^1_n = 1) \\ &+ \mathbb{P}(Z^2_{n-1} = k - 1, Z^1_{n-1} = k - 1, A^1_n = 1). \end{aligned} \tag{61}$$

Now, note that:

$$\mathbb{P}(Z^2_{n-1} = k-1, Z^1_{n-1} = k-1, A^1_n = 1) < \mathbb{P}(Z^1_{n-1} = k-1, A^1_n = 1),$$

$$\begin{aligned}&\mathbb{P}(Z^2_{n-1} = k, Z^1_{n-1} = k)\\ &+ \mathbb{P}(Z^2_{n-1} = k-1, Z^1_{n-1} = k, A^1_n = 1) < \mathbb{P}(Z^1_{n-1} = k),\end{aligned}$$

$$\begin{aligned}&\mathbb{P}(Z^2_n = k) <\\ &\mathbb{P}(Z^1_{n-1} = k) + \mathbb{P}(Z^1_{n-1} = k-1, A^1_n = 1) = \mathbb{P}(Z^1_n = k).\end{aligned}$$

These imply $\mathbb{P}(Z''_n = k) < \mathbb{P}(Z^1_n = k)$, i.e., $v^{*''}_{0,n} < v^{*'}_{0,n}$. □

Appendix F
Proof of Lemma 4

First, equations (26) and (28) can be easily seen by using parts 1) and 2) of Lemma 3 on the welfare formula in (18). Second, equations (27) and (29) can be shown by simplifying the expression of $u^*_{0,n-1}$ for both $\epsilon \in \mathcal{I}_2$ and $\epsilon \in \mathcal{I}_3$, respectively, then substituting into equation (25).

1) Applying Lemma 2 for $k = 2$ and $m \geq 1$, we get:

$$\begin{cases} u_{0,2m+1} = 0 \\ u_{0,2m} = \frac{1}{2} 2^{2m} a^{m-1} (1-a)^{m+1} \cos^{2m-1}\left(\frac{\pi}{4}\right) \sin\left(\frac{\pi}{4}\right), \end{cases} \tag{62}$$

which can also be viewed as a geometric series:

$$\begin{cases} u_{0,2} = (1-a)^2 \\ u_{0,2m+2}/u_{0,2m} = 2a(1-a) \end{cases} \tag{63}$$

This means $u_{0,2m} = (1-a)^2 \left[2a(1-a)\right]^{m-1}$, and:

$$u^*_{0,2m+1} = u^*_{0,2m} = \sum_{\substack{j=2 \\ j\, even}}^{2m} u_{0,j} = (1-a)^2 \frac{1 - [2a(1-a)]^m}{1 - 2a(1-a)} \tag{64}$$

Now, applying (64) to (25), we obtain (27).

2) Similar analysis applies for $k = 3$ and $m \geq 2$. □

Appendix G

Proof of Theorem 4

Refer to Fig. 8. To find τ_n for each agent n, we need to solve, for $\tau_n \in \mathcal{I}_2$, the following equation: $E[\pi_n]|_{\epsilon=\tau_n} = E[\pi_n]|_{\epsilon=\epsilon_3^*}$.

Clearly, we need to consider two cases: $n = 2m$ and $n = 2m+1$. Using the results from (27) and (29), we solve the following equations:

a) For $n = 2m$: solve $E[\pi_{2m}]|_{\epsilon=\tau_{2m}} = E[\pi_{2m}]|_{\epsilon=\epsilon_3^*}$, i.e.,

$$[a(1-a)-\epsilon]\frac{1-[2a(1-a)]^{m-1}}{1-2a(1-a)} = (\epsilon_3^*)^2\,\frac{1-\left(3\epsilon_3^*\right)^{m-1}}{1-3\epsilon_3^*} \tag{65}$$

b) For $n = 2m+1$: solve $E[\pi_{2m+1}]|_{\epsilon=\tau_{2m}} = E[\pi_{2m+1}]|_{\epsilon=\epsilon_3^*}$, i.e.,

$$[a(1-a)-\epsilon]\frac{1-[2a(1-a)]^{m}}{1-2a(1-a)} = (\epsilon_3^*)^2\,\frac{1-\left(3\epsilon_3^*\right)^{m-1}}{1-3\epsilon_3^*} \tag{66}$$

The detailed proof is continued as follows:

i) Showing that $\{\tau_{2m}\}$ is decreasing:

Consider τ_{2m} and τ_{2m+2}. Using (65), to show point A is higher than point C, we need to show:

$$\frac{1-[2a_{2m}(1-a_{2m})]^{m}}{1-[2a_{2m}(1-a_{2m})]^{m-1}} < \frac{1-\left(3\epsilon_3^*\right)^{m}}{1-\left(3\epsilon_3^*\right)^{m-1}}, \tag{67}$$

where $a_{2m} = f(p, \tau_{2m})$.

Let $g(x) = \frac{1-x^m}{1-x^{m-1}}$, for $0 < x < 1$. Then $g(x)$ is increasing. Now note that $\epsilon_3^* = a_3^*(1-a_3^*) > a_{2m}(1-a_{2m}) \Rightarrow 3\epsilon_3^* > 2a_{2m}(1-a_{2m})$, thus (67) is proved, i.e., $\tau_{2m} > \tau_{2m+2}$.

ii) Showing that $\{\tau_{2m+1}\}$ is decreasing:

Consider τ_{2m+1} and τ_{2m+3}. Using (66), to show point B is higher than point D, we need to show:

$$\frac{1-[2a_{2m+3}(1-a_{2m+3})]^{m+1}}{1-[2a_{2m+3}(1-a_{2m+3})]^{m}} < \frac{1-\left(3\epsilon_3^*\right)^{m}}{1-\left(3\epsilon_3^*\right)^{m-1}}, \tag{68}$$

where $a_{2m+3} = f(p, \tau_{2m+3})$. Indeed, we have:

$$\frac{1-\left(3\epsilon_3^*\right)^{m+1}}{1-\left(3\epsilon_3^*\right)^{m}} < \frac{1-\left(3\epsilon_3^*\right)^{m}}{1-\left(3\epsilon_3^*\right)^{m-1}} \tag{69}$$

which is true since $0 < \epsilon_3^* < 1$. Moreover, RHS of (69) is larger than RHS of (68) by applying $g(x)$ on $3\epsilon_3^* > 2a_{2m+3}(1 - a_{2m+3})$. Thus (68) is proved, i.e., $\tau_{2m+1} > \tau_{2m+3}$.

Now, to complete part 1) of Theorem 4, note that both sub-sequences $\{\tau_{2m}\}$ and $\{\tau_{2m+1}\}$ are bounded in $\mathcal{I}_2$, thus they both have limits. Letting $m \to \infty$ in both (65) and (66), these limits are equal the solution of the equation: $\frac{a(1-a)-\epsilon}{1-2a(1-a)} = \frac{(\epsilon_3^*)^2}{1-3\epsilon_3^*}$, which is τ_∞ as in part 2) of Theorem 3, for the case $k = 2$.

2) To show that $\tau_{2m+1} > \tau_{2m}$, we need to show that point C is higher than point D in Fig. 8. This is true because:

$$\begin{aligned}
&[a_{2m+1}(1 - a_{2m+1}) - \tau_{2m+1}]\frac{1 - [2a_{2m+1}(1 - a_{2m+1})]^m}{1 - 2a_{2m+1}(1 - a_{2m+1})} \\
&= (\epsilon_3^*)^2 \, \frac{1 - \left(3\epsilon_3^*\right)^{m-1}}{1 - 3\epsilon_3^*} \\
&= [a_{2m}(1 - a_{2m}) - \tau_{2m}]\frac{1 - [2a_{2m}(1 - a_{2m})]^{m-1}}{1 - 2a_{2m}(1 - a_{2m})} \\
&< [a_{2m}(1 - a_{2m}) - \tau_{2m}]\frac{1 - [2a_{2m}(1 - a_{2m})]^m}{1 - 2a_{2m}(1 - a_{2m})}
\end{aligned} \tag{70}$$

Moreover, $[a(1 - a) - \epsilon]\frac{1-[2a(1-a)]^{m-1}}{1-2a(1-a)}$ is a decreasing function in ϵ, thus (70) implies $\tau_{2m+1} > \tau_{2m}$. $\square$